

**THE RETAILER NEEDS TO BE PAID FOR SOPHISTICATED
DECISIONS: MODELING PROMOTIONAL INTERACTIONS BETWEEN
CONSUMERS, RETAILERS AND BRAND MANAGERS**

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

March 2008

Acknowledgements

The authors acknowledge financial support from 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 by choosing the store promotions that consumers see each week. The interactions between consumers, retailers and brand managers are complex, with each having different goals and constraints. Here we use regression and Genetic Algorithm techniques to model these interactions. Our goal is to address an important question in the literature; specifically whether the retailer gains positive returns through more sophisticated moderation. In doing so, we draw on work in economics showing that unsophisticated or Zero-Intelligence (ZI) agents can often replicate real-world outcomes. Hence we posit and test three levels of retailer sophistication—ZI, naïve and sophisticated—bringing a novel perspective to this important topic.

We find that, although the behavior of brand managers 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 that the retailer either receives adequate side-payments from the dominant brand or there are sufficient spin-off benefits in other categories to compensate the retailer for the loss of volume in the focal category.

Keywords: store promotions, retailer behavior, Genetic Algorithm, zero-intelligence agents

1. 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 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 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 seller that improves their brand's weekly profits compared with historical results for an actual market. Here 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 their marketing instruments. Our approach has been to use econometric models of consumer response to marketing instruments, bit-string agents modeling the decisions of brand managers and the Genetic Algorithm (GA) to adapt these agents to consumer response. 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.

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 store promotion in a category in any week and only for one week.

In 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, 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. 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 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.

There are two streams of marketing literature that address the role of the retailer and which shed light on how they might go about choosing store promotions. One literature is primarily empirical, the other more analytical and game-theoretic. But these literatures present something of a paradox. The empirical literature suggests that the retailer's choice is made simply, often only looking at the focal brand and two or three promotional variables. In contrast the analytical literature suggests attention to greater numbers of brands and variables—as in category management—would produce benefits. Furthermore, the empirical literature suggests that store promotions are often disadvantageous to the retailer but the two literatures disagree on whether side payments can compensate for this. The empirical literature finds little scope for side payments but the analytical literature argues they are optimal. Thus incorporating the retailer's choice of promotions into our models might also help resolve this paradox, especially as our approach incorporates both empirical and analytical aspects. But the question remains how sophisticated need this choice be?

Here an important contribution from economics needs to be incorporated into the store promotion literature. Gode & Sunder (1993) discovered that a continuous double auction market of trading agents with virtually *no rationality* still approximates the equilibrium price. Since their discovery there has been an increasing awareness that allocative efficiency might well be an emergent characteristic of the market, rather than the result of sophisticated decision-making by agents. Researchers have begun to explore

what elements of a market are necessary for this emergence. Such explorations would not have been possible without the tool of computer simulation, to conduct what-if experiments, and the use of so-called Zero Intelligence agents, thus separating the market from the unobservable goals of human traders.

Following these ideas, we explore the moderation of an iterated oligopoly of coffee brands vying for revenues and profits in the supermarket. The moderation occurs because the supermarket is not a neutral party, but wishes to harness the brand rivalry to improve its profits too. In this paper we posit and test three levels of moderation by the retailer: (1) zero-intelligence 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 retail model before. 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 proposed in any week, the type of promotion and the proportion of stores in the chain that will run the promotion. 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 brand asymmetries in 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 review the marketing literature on promotions and side payments and also our own prior work. Then we show how we develop and implement the naïve and sophisticated retailer models, outline the literature on ZI agents and formulate our hypotheses. Then, in the methodology section, we present our Monte Carlo approach, the range of experiments that we ran, and show how we analyze the results of these experiments. Finally, we present the results of these analyses, discuss these in more depth, and close with our conclusions and suggestions for further work.

2. Prior Work

A. The Store Promotion Literature

Store promotions, particularly in supermarkets, are a major topic in the marketing literature. But it is fair to say that historically more attention has been paid to their effects on consumer demand (brand switching, purchase acceleration, etc.) than to supply-side issues, especially those concerning retailers. For example, using the same product category as does this paper, Gupta (1988) finds that 84% of the

impact of a price promotion is due to brand switching and 16% to purchase acceleration. Other authors such as Bucklin, Gupta and Siddarth (1998), Bell, Chiang and Padmanabhan (1999) and Pauwels, Hanssens and Siddarth (2002) have replicated these findings across multiple product categories, although noting that there are large differences between categories. Bell, Chiang and Padmanabhan show that these differences relate systematically to category-specific factors such as necessity, share of consumer budget, brand assortment, etc. Neslin (2001) provides a review of this literature.

More recently, the focus has switched, in part, to the short- and long-term effects of promotions. Authors such as Nijs, Dekimpe, Steenkamp and Hanssens (2001) and Pauwels, Hanssens and Siddarth (2002) have shown that, in general, promotions neither generate long-term benefits to the promoting brand nor do it any permanent damage. But here too most studies look at *consumer demand* and not brand or retailer *revenue* or *profits*. And few examine cross-category effects between complements or substitutes for the focal category (Kamakura and Kang 2007). Thus empirical knowledge on the economics of promotions is largely limited to inferences about volume effects in single categories.

Srinivasan, Pauwels, Hanssens and Dekimpe (2004) take up this challenge by asking whether store promotions benefit manufacturers, retailers or both. Their finding of relevance here is that, in general, promotions are more attractive in *revenue* to manufacturers than to retailers, possibly because the retailers' revenue loss on non-promoted items is the same or more than their gain on the promoted item.¹ They also investigate whether these category losses were compensated by gains in *overall* store traffic and revenue, finding that only some *national* brands generate positive gains for the retailer. Finally, they look at the scope for compensatory side payments from manufacturer to the retailer, finding little scope even amongst these national brands.²

¹ Another possible explanation they advance is that prices stay below their baseline for some weeks after the promotion. We did not observe this in our data which is for a different chain, with different policies and managerial controls.

² Note they, like us, have no knowledge of actual side payments (for example, slotting fees paid by manufacturers). Their conclusions are made by comparing estimated revenues.

Expanding on the theme of who benefits, Besanko, Dube and Gupta (2005) look at the extent to which the retailer passes the manufacturer's promotional discount through to the consumer. The rate of pass-through is a critical decision for the retailer, trading off margin improvement against the impact of the promotion on consumers. Besanko *et al.* find that the more popular the brand is to consumers, or the more profitable the brand to the retailer, the more likely the retailer to pass the manufacturer's discount through. They suggest two possible reasons for this.

First, promotions on popular brands may expand overall category sales rather than cannibalize other brands in that category. This suggestion is consistent with Bronnenberg and Mahajan (2001), who argue the retailer is more likely to pursue volume objectives in large categories or for national brands. Promotions in important categories or for popular brands can draw attention to the retailer and increase store traffic. Retailers may even cross-subsidize such activities by so-called "loss leaders."

Second, for popular or expensive brands the retailer is less likely to adjust the prices of other brands during the promotion. In contrast, smaller share brands suffer a triple whammy. They receive lower pass-through, the retailer reduces the prices of popular brands during their promotions, and their own retail prices do not fall to compete during promotions for popular brands.

Pauwels (2007) also looks at pass-through and the extent to which the retailer adjusts the prices of competing brands during promotions. Across 25 categories Pauwels finds that the average pass-through rate is 65%. This average rate implies a long-term wholesale promotional elasticity of 1.8, but this reduces to 1.6 because competitors typically drop their wholesale prices by 15% during promotions.

In Pauwels' data, pass-through rates vary by category size, brand share, how expensive the brand is, and category concentration. Larger categories and higher brand shares invoke greater retailer pass-through. As above, Pauwels argues this may be to achieve traffic-building objectives. Pauwels reaches the overall conclusion that brand managers typically perceive minimal damage from the promotions of their competitors, due again to category expansion effects.

The empirical literature is relatively silent on retailer policy towards promotions. Anecdotal observation suggests some retailers allow competing promotions, whereas others, including the focal

retailer in our data, only allow one brand to be on promotion in any one week. This does not preclude other brands offering discounts to compensate but it does reduce the magnitude of competitive effects, especially as non-promoted brands are less able to draw the attention of consumers to any lower prices that the retailer passes through. Similarly, intuitively, it seems less likely that a retailer who has the policy to highlight one promotion would compromise this policy with less noticeable price reductions on other brands, an intuition that has some support in the category management literature. Category management is a methodology where retailers jointly consider demand, costs, and prices of all the brands in the category when setting prices or choosing promotions for a focal brand (Zenor 1994). Such considerations can lead retailers to prefer to promote only one brand at a time (Leeflang and Wittink 1992, Tellis and Zufyrden 1995). The downside is that category management implies more complex decision-making, which has costs for the retailer.

The issue of what drives retailers to set prices or choose promotions is of course potentially broader than just the category itself. Nijs, Srinivasan and Pauwels (2007) examine all the various drivers of retail pricing through two data sets covering 67 categories of product. These drivers include the prices of competing retailers, thus acknowledging that the retailer itself operates in a market. But, in common with earlier studies, Nijs et al. find competition has only a small influence on the prices set by the retailer. Category management has a similarly small influence. Nijs *et al.* conclude that when setting prices retailers look at (1) past prices, (2) likely demand for the brand, and (3) wholesale prices. In contrast to category management or considerations of retail competition, these results suggest relatively simple decision-making by retailers.

These conclusions gain further support when these authors look at pricing and retail margins. Nijs et al. find that when retailers over-emphasize traffic building or the importance of past prices, margins are lower. When retailers emphasize demand or category management considerations, margins improve. Nijs et al. interpret this result as implying that a simple cost-plus pricing approach neither hurts nor benefits the retailer, whereas a more sophisticated approach might benefit the retailer. But unknown in all these studies are the costs of sophistication to the retailer, including information, analysis and time.

Taken one category at a time these may not be that significant, but retailers have to manage a large number of categories (30,000 SKUs in a typical supermarket). Many interesting research questions remain about the extent to which sophisticated decision-making is beneficial, for which categories and under what circumstances. These are important questions because, whereas the marketing literature universally assumes sophistication is beneficial, the ZI literature shows that in some circumstances simplistic decision-making can bring almost the same benefits.

A second stream of marketing literature looks at all these issues from an analytical modeling perspective. Gertsner and Hess (1995) show that promotions help alleviate the double marginalization problem in channels and argue the right combination of store promotions and side payments should increase total channel profit. Kim and Staelin (1999) build a one-period profit-maximizing game around a channel containing two manufacturers and two retailers. Here both manufacturers and retailers face competition, and their results show that in such circumstances it is in the manufacturers' best interest to offer side payments and it is in the retailers' best interest to pass some of these payments through to consumers. Similarly Silva-Risso, Bucklin and Morrison (1999) use an empirically calibrated simulator to show that side payments are optimal for the manufacturer even when this pass-through is low. Ailawedi (2001) reviews this and other literature in the context of retailer-manufacturer power and performance. He notes that important questions remain unanswered about the many forms of side payments that are observed, including their impact on the profitability of manufacturers and retailers.

More recently, Moorthy (2005) uses comparative statics to examine pass-through as a more general phenomenon of prices and costs. His work shows that cost changes on one brand can have positive or negative effects on the prices of others. Bensanko *et al* (2005) note there are a wide range of possible cross-brand pass-through effects; depending on the assumptions made about the behavior of the retailer. For example, a vertical Nash Equilibrium model yields zero cross-brand pass-through. In contrast, a category management model (with certain demand effects) predicts that the prices of non-focal brands will go up, whereas if the retailer acts as a brand manager, the prices of non-focal brands will go down.

Raju and Zhang (2005) consider how best the manufacturer can coordinate a channel in the presence of a dominant retailer. They show that both quantity discounts and two-part tariffs can achieve this coordination, although not equally efficiently. Raju and Zhang also show that “street money” (slotting allowances) can arise as the minimum incentive to a dominant retailer to coordinate with the manufacturer. Kuksov and Pazgal (2007) also look at slotting allowances in the context of two-part tariffs between manufacturers and retailers. They find that more intense retail competition, higher retail bargaining power, larger retail fixed costs, lower marginal costs of retailing, and larger relative retailer size all encourage the manufacture to offer slotting allowances.

These two streams of literature present something of a paradox. From the first literature, store promotions appear financially disadvantageous to the retailer in revenue and margin terms, except perhaps for promotions involving some national brands. Nor do the revenue and margin estimates leave much scope for side payments. From the second literature, there appear to be good theoretical arguments that side payments are optimal, which can also be coupled with the well-known fact that their use and sophistication is increasing (e.g. Dreze and Bell 2000).

Equally, the second literature suggests that, from a theoretical perspective, optimizing profits requires the retailer to consider many demand- and supply-side variables, including some not found in empirical data sets. Yet, the first literature finds actual retail decision-making to be fairly simple, often focusing narrowly on the focal brand and three variables (past price, wholesale prices including trade deals, and likely brand demand). Is this because of a lack of managerial sophistication in pricing, as some authors suggest (for example, Krishna, Mela and Urbany 2001), or because, as the ZI literature indicates, greater sophistication is not really necessary?

B. Our Previous Model

In Midgley, Marks and Cooper (1997) we showed how to 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.

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 manageable, we use cluster analytical techniques to define a smaller set of typical actions. This becomes the response menu from which the bit-string agent representing each brand can 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 a response from this menu that will be profitable for it in the next period of the game, given the stimulus of the previous state of the market. The GA is used to find mappings between stimulus and response that maximize profits over many multi-period games, pitting various agents against their competitors. The GA does 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.

3. 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. The

more interesting case is if two or more brands offer a promotion for the same period. In our prior work the retailer follows the ZI approach and simply tosses a coin to decide which brand to promote. In doing so the retailer takes no account of the features of the competing promotions. This paper extends the retailer decision-making to the naïve and sophisticated levels mentioned above, 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, made available 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 price often. Although they are present in the consumer response model, we do not include them in the retailer models.

We choose *total category volume* (pounds per week) as the dependent variable. This variable is reasonable as a proxy both for retailer profit (which we do not have) and for the impact on sales of other categories in the store, a position that has support in the literature (e.g. Bronnemberg and Mahajan, 2005; Pauwels, 2007). Total category volume is thus a useful criterion for deciding between two competing promotions. The retailer should prefer the promotion with the greatest positive impact on weekly sales.

The independent variables in the data include promotional and shelf prices and the proportion of stores 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: Our naïve retailers are 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 implement 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-square of 65% and the coefficients shown in the upper part of 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.

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, using standard autoregressive procedures. Exploratory analyses also revealed that we should include the impact of sales at a competing chain.

Several alternative model formulations were tried; including models where we transform the independent variables to lessen multicollinearity by weighting the discount by the proportion of stores in which the promotion was displayed, and models with these display and discount variables kept separate. Some of the latter gave the best R-squares, but also exhibited greater multicollinearity, mainly because depth of discount correlates highly with the display proportion. Thus for these models the *incremental effect* of promotions is less well calibrated (larger standard errors on coefficient values) than for the weighted discount models. As this incremental effect is our main focus, we prefer the latter. The final model is shown in the lower part of Table 1, with an R-squared of 88% and 66 degrees of freedom.

This model represents a sophisticated retailer. Such retailers 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 other chains 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 1 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 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 2.

[Table 2 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), with store displays and store advertising. The other is a minor promotion, with a smaller discount and only store displays. In our data these occur with a frequency of two major to one minor. As we have three actions remaining of our eight available, 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. Some outliers were excluded for minor promotions. We believe these to be mistakes; basically promotions that store managers forgot to remove

at the start of the following week. Otherwise, this procedure resulted in two homogenous clusters for major promotions (shown as LOW1 and LOW2 in Table 2, and one for minor (shown as 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, as we did 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 still 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 3 shows the predicted volumes that each promotion will achieve (ranked from highest to lowest impact).

[Table 3 about here]

Taking the naïve retailer first; any promotion results in significantly increased volume. Indeed, all volumes are about 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. Thus our retailers are also in some ways category managers who maximize the impact of the promotion on category volume by keeping all other brands at their shelf-price.

E. 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. They 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. 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.

4. 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 ($4,096 \times 3$). 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. To make this representation work, we use the GA to search for good state-response mappings. Specific details are available from the authors.

B. Using the GA to optimize these strings

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 12,300 bits long, with four brand populations..

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 consider other strategies that may exist in real markets, such as cross-subsidization across the various brands owned by one manufacturer, *et cetera*.

Each agent-string plays a 50-round game with every possible combination of strings from the other three brands' string populations. Its average profit is its fitness. Testing each generation required 2,925 games per agent-string, or 73,125 50-round games. This is a non-trivial simulation, especially as it takes around 100 generations to achieve satisfactory convergence for the four brands, which equates to 29,250,000 50-round games.

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. 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 are probably 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 Nash Equilibria – the combinations of best responses, where "best" might be constrained by the bounded rationality of the human manager.

Can we seriously claim to search this vast space of 25 agents per brand 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. Grefenstette (1986) argued 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 an agent population of 25. This can be seen from our 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 across 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 multiplied by 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 analyze the historical actions of brands and retailers. We do this by clustering shelf prices using the stylized actions shown in Table 3 as cluster seeds which, combined with the promotional analysis 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.

5. Results

A. Actions of the brand-agents

Table 4 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 4 is that when a sophisticated retailer moderates promotions, the outcome is much closer to the historical pattern, indeed the closest we have obtained in our program of research. Thus the form of retailer has a major role in controlling the promotional behavior of the brands.

[Table 4 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 5 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 ($p < 0.002$), although 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. But 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.

[Table 5 about here]

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 6 shows the simulation results.

We omit the ZI case, where 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, as 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 naïve retailer is blind to brand. The other notable feature of the naïve retailer clashes is that nearly every week of the 50 weeks involves a clash. Only twice did a brand propose a promotion without competition.

[Table 6 about here]

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 naïve 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 6. What are the reasons for these differences between the naïve and sophisticated cases?

Each brand has three promotional actions (LOW1, LOW2, and Low1 in Table 2). 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 naïve retailer, the descending ranking is simply: LOW1, LOW2, and Low1, since the naïve 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 3. 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 6 are the realized frequencies of clashes (and outcomes) which occurred in the experiments.

7. 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 5). 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 4). Finally, with sophisticated moderation, when the retailer looks at which brand is proposing a promotion, as well as the promotional market actions (Table 3), the brand-agents learn (in an evolutionary sense) to avoid promotional clashes (Table 6).

Should we conclude that it is difficult to justify the greater costs (whatever they are) of sophisticated moderation by improved category volume? It is not true that naive moderation strictly dominates sophisticated moderation, since *Folgers* preserves its profits with sophistication, and *Folgers* is the dominant brand. *Folgers* 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 unknown to us as we do not have that data. It is also possible the spill-over effects of promoting the dominant brand on sales in other product categories compensate the sophisticated retailer, but again we lack data. This could be a reasonable explanation, though, as few consumers shop for just one item. And of course, this moderation reduces the profits of *Folgers*' 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 our agents learn 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. Limitations

The obvious limitations to our work are first that it only looks at one regional market for one product category. Replication is needed to generalize these findings. Equally we only consider three of many possible models of retailer choice, which in two cases we assume to follow linear formulations. Other types of decision making processes may be equally or more valid. Finally, we are limited in the length of our data series, preventing consideration of short- versus longer-term impacts of promotions.

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. 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, needs either more extensive datasets or simpler, perhaps stylized, models of other categories from more limited data.

Finally, and much more ambitiously, we can replace both econometric models of 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.

References

- Ailawedi, K.L. 2001. The retail power-performance conundrum: What have we learned? *Journal of Retailing*, **77** 299-318.
- Art, D., R. Gnanadesikan, R. Kettenring. 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. L. Davis, ed. *Genetic Algorithms and Simulated Annealing*, Pittman, London, 32-41.
- Bell, D., J. Chiang, P. Padmanabhan. 1999. The decomposition of promotional response: an empirical generalization. *Marketing Science* **18**(4) 504-526.
- Besanko, D., J.P. Dubé, S. Gupta. 2005. Own-brand and cross-brand retail pass-through. *Marketing Science* **24**(1) Winter 123-137.
- Bronnenberg, B.J., V. Mahajan. 2001. Unobserved retailer behavior in multimarket data: joint spatial dependence in market shares and promotion variables." *Marketing Science*, **20**(3) 284-299.
- Bucklin, R. E., Gupta, S., S. Siddarth. 1998. Determining segmentation in sales response across consumer purchase behaviors. *Journal of Marketing Research*, **35**(2) 189-197.
- Carpenter, G., L. Cooper, D. Hanssens, D.F. 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, November.
- Cooper, L., M. Nakanishi. 1993. *Market share analysis*. Boston, MA, Kluwer.
- De Jong, K.A. 1975. An analysis of the behavior of a class of genetic adaptive systems. PhD dissertation.

- Dreze, X., D.R. Bell. 2003. Creating win-win trade promotions: theory and empirical analysis of scan-back trade deals. *Marketing Science* **22**(1) Winter 16-39.
- Duffy, J. 2006. Agent-based models and human subject experiments. The *Handbook of Computational Economics, Volume 2: Agent-Based Modeling*, edited by L. Tesfatsion and K.L. Judd. The series *Handbooks in Economic*. General editors K. Arrow and M. D. Intriligator, Amsterdam, Elsevier Science, 949-1011.
- Gathercole, C., 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, 111-118.
- Gertsner, E. J.D. Hess. 1995. Pull promotions and channel coordination. *Marketing Science* **14**(1) Winter 43-60.
- Gode, D., 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) January-February 122-128.
- Gupta, S. 1988. Impact of sales promotion on when, what, and how much to buy. *Journal of Marketing Research*, **25**(4) 342-356.
- Kamakura, W.A., W. Kang. 2007. Chain-wide and store-level analysis for cross-category management, *Journal of Retailing* **83**(2) 159-170.
- Kim, S.Y., R. Staelin. 1999. Manufacturer allowances and retailer pass-through rates in a competitive environment. *Marketing Science* **18**(1) 59-76.
- Krishna, Aradhna, C.F. Mela, J. Urbany. 2001. Inertia in pricing. Working paper, University of Notre Dame, Notre Dame, IN.
- Kuksov, D., A. Pazgal. 2007. Research note: the effects of costs and competition on slotting allowances. *Marketing Science* **26**(2) March-April 259-267.

- Leeflang, P.S.H., D.R. Wittink. 1992. Diagnosing competitive reactions using (aggregated) scanner data. *International Journal of Research in Marketing* **9**(1) 39-57.
- Lu, Q., S. Moorthy. 2007. Coupons versus rebates. *Marketing Science* **26**(1) January-February 67-82.
- 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., D.F. Midgley, L.G. Cooper. 2006. Co-evolving better strategies in oligopolistic price wars. *Handbook of Research on Nature-Inspired Computing for Economy and Management*. Edited by Jean-Philippe Rennard, Hershey, PA: Idea Group Inc., 806-821.
- Midgley, D.F., R.E. Marks, L.G. Cooper. 1997. Breeding competitive strategies, *Management Science* **43**(3) March 257-275.
- Moorthy, S. 2005. A General theory of pass-through in channels with category management and retail competition. *Marketing Science* **24**(1) Winter 110-122.
- Neslin, S. A. (2001). Sales promotion. *Handbook of Marketing*, eds. Barton Weitz and Robin Wensley, Sage Publications.
- Nijs, Vincent, M.G. Dekimpe, J-B.E.M. Steenkamp, Dominique M. Hanssens. 2001. The category demand effects of price promotions. *Marketing Science* **20**(1) 1-22.
- Nijs, V.R., S. Srinivasan, K. Pauwels. 2007. Retail-Price Drivers and Retailer Profits. *Marketing Science* **26**(4) July-August, 473-487.
- Pauwels, K., D.M. Hanssens, S. Siddarth. 2002. The Long-term effects of price promotions on category incidence, brand choice and purchase quantity. *Journal of Marketing Research* **39**(4) November 421-139.
- Pauwels, K. 2007. How retailer and competitor decisions drive the long-term effectiveness of manufacturer promotions for fast moving consumer goods. *Journal of Retailing* **83** (3) 297-308.
- Raju, J., Z.J. Zhang. 2005. Channel coordination in the presence of a dominant retailer. *Marketing Science* **24**(2) Spring 254-262.

- Schraudolf, N.N., J.J. Grefenstette. *A User's Guide to GAucsd 1.4*, Technical Report CS92-249, CSE Department, UC San Diego, 1992, <http://www.agsm.edu.au/~bobm/teaching/SimSS/GAucsd.pdf>
- Silva-Risso, J.M., R.E. Bucklin, D.C. Morrison. 1999. A decision support system for planning manufacturers' sales promotion calendars. *Marketing Science*, **18**(3) 274-300.
- Srinivasan, S., K. Pauwels, D.M. Hanssens, M.G. Dekimpe. 2004. Do promotions benefit manufacturers, retailers, or both? *Management Science* **50**(5) May 617-629.
- Tellis, G.J., F.S. Zufryden. 1995. Tackling the retailer decision maze: which brands to discount, how much, when and why? *Marketing Science* **14**(3) 271-299, Part 1 of 2.
- 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.
- Zenor, M. J. 1994. The profit benefits of category management. *Journal of Marketing Research*, **31**(2) 202-213.

TABLE 1: Regression coefficients for Retailers

Variable	Coefficient	t-value	p value
Naïve Retailer (Simple OLS model)			
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
Sophisticated Retailer (Weighted discount form)			
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
Major Promotions			
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
Minor Promotions			
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
Impact of other supermarket chain			
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

Note: Shaded coefficients are those relevant to a choice between competing promotions

TABLE 2: Stylized actions

Action label	Price (\$)	Feature (% of stores)	Display (% of stores)
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 3: 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 4: Stylized actions from Monte Carlo simulations (Percentage of actions)

Stylized Actions

	LOW1	LOW2	Low1	Low2	High1	High2	HIGH1	HIGH2
Zero-Intelligence Retailer								
<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
Naive Retailer								
<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
Sophisticated Retailer								
<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
Historical								
4 brand data	2	7	2	13	24	12	33	7

TABLE 5: Total category volumes and brand profits

	Zero-intelligence retailer		Naive retailer		Sophisticated retailer	
	Mean	Std.dev	Mean	Std.dev	Mean	Std.dev
Total Category Volume	1756	17	1763	40	1617	63
Profits						
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 6: 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				