Artificially Intelligent Specification and Analysis of Context-Dependent Attribute Preferences

Bruce Edmonds, Scott Moss and Helen Gaylard
Centre for Policy Modelling
Manchester Metropolitan University
Manchester M1 3GH
http://www.cpm.mmu.ac.uk/

Abstract

In many markets, customer preferences are context dependent. In the professional marketing literature, this dependence is typically recognized as a “need-state”.Moss and Edmonds (1997) recently reported a model that allows the testing of the qualitative judgements of domain experts in spirits markets against relevant EPOS data of product sales. This paper extends the use of context dependent customer preferences to the case where domain experts lack confidence in their judgements or the judgements are found not to be in accordance with the data. We describe here an algorithm to produce credible alternative models for the domain experts to confirm or develop in light of their wider domain expertise.

The algorithm combines random search, genetic programming and evolutionary hill climbing techniques. We report the results of tests using data from markets for alcoholic beverage. The algorithm enabled the largely endogenous production of qualitative descriptions which are both consistent with observed data and deemed credible by domain experts. In a detailed example, the technique is shown to provide extensive insights into the reason for a widely successful brand to have made little impact in one geographically defined market.

The algorithm and its implementation are as rigorous and accurate as conventional, purely statistical techniques. They have the additional advantage of cohering with the language of discourse of the domain experts.

Keywords: brand choice, choice models, market structure, buyer behaviour, artificial intelligence, econometric modelling

1 Introduction

Applications of utility theory represent each consumer as a preference function which is independent of the reason for which any purchase is to be made. Moss and Edmonds = [23], however, have reported an intelligent marketing integrated system (IMIS) incorporating a view from the marketing profession that, in some markets at least, preferences are usefully related to the context of consumption. They demonstrated a specification of context-dependent attribute preference (CDAP) functions that tracked the relationship between prices and market shares of branded fast-moving consumer goods (FMCGs) better than corresponding OLS models when estimation was based on the short sample periods. That is, the CDAP model incorporated the qualitative judgements of domain experts but used less statistical data to produce more accurate market share simulations than did the OLS models. It also provided substantial qualitative information that was useful in determining marketing strategies. Such qualitative information cannot, of course, be provided by statistical models of any kind.

Moss and Edmonds did not argue that CDAP models are in some sense better than statistical models. They have different purposes and, in any case, integrating statistical analysis with CDAP analysis leads to a better grounded and more fully rounded picture of competitive relations in some markets than does either approach by itself.

CDAP models incorporate of the qualitative representations of demand factors in a way which enables us to test their consistency with the available statistical record. We follow the formal modelling literature in which "qualitative" relations are functional numerical relations which are invariant under a set of topological transformations. See, for example, [15, 9, 5, 26, 11]. In particular, the qualitative relations used here are mappings from verbal expressions into real number intervals.

Clearly, this approach is very different from the concerns of the literature on consumer behaviour typified by [1] or [8]. It is not concerned directly with how individuals formulate their preferences or their purchasing and consumption decisions. Instead, preferences are represented as a distribution of the attributes required by consumers who are purchasing branded goods for specific purposes such as alcoholic beverages for parties or for consumption individually within the home.

For each purchasing context, the distribution of attribute ideals reflects differences in individual tastes and also the tolerance of consumers to differences between the ideal intensity of an attribute and the actual value which, in their perception, is embodied in actual brands. Moreover, particular attribute values will be more or less critical to consumers purchasing for different purposes. Consequently, for each purchasing context identified by the marketing professionals, we identify the attributes which are identified with the various brands as well as the ordering of the intensity of such attributes. Thus, Chivas Regal whisky has more of the attribute "specialness" than does Harry's Old Horsegut Bourbon but Harry's Bourbon has more of the attribute "special ingredients" than does Chivas Regal. When celebrating a successful event, consumers might give a high priority to buying a drink which has a high degree of "specialness" but when contemplating the rich variety of life they might want something more distinctive. In the one case specialness will be important and in the other case it might not. The contemplatives, taken as a group, might have widely differing ideal levels of distinctiveness in their drink or they might be tolerant of substantial deviations from their ideal of the perceived uniqueness of a drink.

A detailed justification for this approach is given in section 2. The preference distribution functions are motivated and described in detail in section 3, the algorithms for extracting credible CDAP models from EPOS data is described in section 4 and then, in section 5, applied to the case of a market for alcoholic beverages.

The results reported in section 5 demonstrate that the formal incorporation of the verbal, qualitative judgements of domain experts into a model of a specific market supports the extraction of detailed qualitative information which is demonstrably consistent with reliable, numerical, time-series data. This result is consistent with earlier findings such as the demonstrations by [4] and by [22] that adding rulebased descriptions of expert knowledge into econometric forecasting models can substantially improve their accuracy. Indeed, there is a substantial and long-standing literature on the incorporation of qualitative, expert knowledge in formal models. See, for example, [21, 18]. There is also the well established literature on the reliability measures for qualitative judgements made independently by several domain experts. [24] review and extend this literature.

2 The competitive set

The data sets used to test the models reported in this paper cover at least 65 and up to nearly 200 brands of alcoholic beverage. The individual brands are types of whisky

(Scotch, Irish, Canadian, Bourbon, etc.), white spirits (gin and various sorts of vodka), fortified wines, brandies and liqueurs. Choosing any one brand in such a data set, it is by no means obvious which of the other brands are in its set of competitors. A Scotch whisky, for example, could compete with brandies, other types of whisky or even liqueurs and fortified wines in different contexts.

Using algorithms driven by a knowledge-based system as reported by [3], we inferred from data on prices and sales volumes for all brands in the data set the chief price-competitors of the brand we chose as the focus of each simulation run with the CDAP model. These algorithms have been developed and integrated *ad hoc* and the demonstration of their formal properties is reserved for further research. Nonetheless, they do guide and inform the specification of the competitive sets by the marketing professionals and, so, we use them here in full recognition of their possible formal weaknesses. Our justification is that the knowledge-base describes the actual (though *ad hoc*) procedures used on actual data sets to inform the development of marketing strategies in earnest.

The determination of the competitive set of a focus brand proceeds in three stages:

1) the determination of a plausible superset of the competitive set by linear regression of market share of the focus brand on some transform of the price of each of the other brands for which data is held as well as regressing the shares of the other brands on the same price variable of the focus brand;[1]
2) the elimination of some brands from that superset by a set of multiple regressions developed from the AIDS algorithm but without the symmetry restriction mentioned above;
3) some further elimination of brands from the competitive set together with analysis of the changes in competitive structures over the data period based on a non-linear (local-regression-based) generalization of the second stage.

Because the number of brands is often large in relation to the number of observations (up to 2500 brands with no more than 200 observations) there are insufficient degrees of freedom to begin with the second or third steps involving multiple regressions over all other brands.

Effectively, the first cut at the competitive set was to include all brands for which all of the relevant regression coefficients were of the correct sign with t-statistics greater than 5 in magnitude. The relevant coefficients were the OLS coefficients on price with focus-brand value share as the dependent variable and, where these were of appropriate sign and magnitude, the coefficient on the focus brand price with the competing brand's value share as the dependent variable. The point here is to ensure that the competition goes both ways even if the cross-price elasticities are not symmetrical.

The second cut used similar criteria but in an OLS regression of the value share of the focus brand against a log transform of all of the price variables in the first-cut competitive set. Brands were winnowed out of this set one at a time either because in the regressions their coefficients had the wrong sign (indicating that they were complements rather than competitors) or, if all coefficients were of the appropriate sign, because their coefficients were the least significant.

---

[1]In all of these cases the log of total volume and time were included as regressors to eliminate the effects of seasonal fluctuations in overall demand as well as demand trends.

At this stage, it is usual for some surprising brands to be left in the competitive set. If the set has been cut down too far, some brands that the domain experts would expect to be in the competitive set are left out. Since we want to be sure that no actual competitors are left out of further consideration, we typically make the second-cut competitive set rather larger that the size of the set we intend to end up with. In general, the marketing professionals are interested in the half-dozen or so most important competitors. If we leave 15 to 20 brands in the second-cut competitive set, then they have some confidence that a set of that size will include all of the most important six to eight competitors.

The reason that inappropriate brands are left in the second-cut competitive set is that some ephemeral strategy has brought them temporarily into competition with the focus brand. Usually, this will be because of some special offer which increases their sales while the offer remains in force but does not lead to a long-term increase in market share. The problem here is that the assumption of a linear relation between market share and competitors' prices may yield a spurious result in which a few large and systematic fluctuations in volumes and prices are averaged out over all observations and make the constant coefficients and t-statistics larger and apparently more significant than would be the case without those few fluctuations.

In order to identify the brands which should be in the competitive set but might not be captured by linear regressions and their interpretations as well as those that should not be captured, we employ non-linear, local regression. This procedure produces a regression coefficient for each observation and each regressor.[2]

The interpretation of the time patterns of local regression coefficients is determined by rules and is entirely declarative. The rules "look" at the patterns of levels and first and second differences in the coefficients on each regressor over the data period. The coefficients of interest are those on the price variables where the dependent variable is the focus brand's value share. These will be significant and positive for competing brands and insignificant or negative for non-competing brands. A positive coefficient indicates that the focus brand will lose (*resp.* gain) share if the price of another brand falls (*resp.* rises). Because the price variable used is a log transform of the actual price, the coefficient is the elasticity of focus-brand value share with respect to the price of the other brand. Thus, a high coefficient value indicates a high elasticity and, therefore, more competitiveness.

The aim of the rulebase in this regard is to identify brands which are consistently strong competitors of the focus brand and to include brands which become competitors over the data period and to eliminate from consideration those which have ceased to be competitors during the data period. The time pattern of the local regression price coefficients identifies which brands fall into these various categories.

3 CDAP functions

Our approach in general is to represent products by their attributes, the intensities of the respective attributes, their market strengths and their prices. These elements of the products determine their market shares.

The underlying paradigm of the model is that the market share of any one product is taken from the shares of other (but possibly not all other) products in its competitive set. A relative price reduction, for example, will take market share from other products and, in particular, those other products which are most similar in terms of perceived attributes. A

[2] A full description is contained in [3].

relative increase in market strength brought about by a successful marketing campaign will similarly take share away from other, in some sense similar, products. We do not assume any symmetry in these relations. A product with much greater market strength will have a greater effect on the share of much weaker products than will the weaker products on the stronger.

Market-level parameters determine the relative importance of product differentiation, price differences and context-dependent preferences in each market. These parameters are determined endogenously because there is no reason to assume that these factors will have the same effects in different markets.

3.1. Describing the market: the reach function

In order to capture these ideas in a model we define a variable which we call *reach*. This variable is an index of the share which one product takes from another. *Reach* is larger the greater the relative market strength and the lower the relative price. But the effect of either market strength or price is less as there is less similarity between the products.

Denote by ρ_{ij} the reach of the *i*th with respect to the *j*th of a set of *n* products. Since we intend to use this concept of reach to determine the volume shares of the various products, all *n* of the products must be similar in the sense that their quantities can be measured in some common unit such as litres or grams or, in the case of non-financial services, person-hours.

Though we will develop a formal measure of market strength presently, we simply assert at this stage that there is some consistent measure of the market strength of each brand and denote by σ_i the market strength of the *i*th brand. The price is p_i. The vector of indices of attribute intensities of the *i*th brand is Θ_i. Just how we obtain these and what they mean will be described in general terms below and by example in the next section. In the usual notation, the distance between two attributes vectors is $|\Theta_j - \Theta_i|$.

Formally,

(1) $$\rho_{ij} = \rho_{ij}\left(\frac{\sigma_i}{\sigma_j}, \frac{p_i}{p_j}\right)$$

where

(2) $$\frac{\partial \rho_{ij}}{\partial\left(\frac{\sigma_i}{\sigma_j}\right)} = \rho_\sigma\left(\frac{\sigma_i}{\sigma_j}, |\Theta_i - \Theta_j|\right) > 0$$

(3) $$\frac{\partial \rho_{ij}}{\partial\left(\frac{p_i}{p_j}\right)} = \rho_p\left(\frac{p_i}{p_j}, |\Theta_i - \Theta_j|\right) < 0$$

(4) $$\frac{\partial \rho_\sigma}{\partial\,(|\Theta_i - \Theta_j|)} < 0$$

(5) $$\frac{\partial \rho_p}{\partial(|\Theta_i - \Theta_j|)} < 0$$

Verbally, the reach of one brand with respect to another is determined by their relative market strengths and relative prices. Naturally, reach increases with relative strength (inequality (2)) and diminishes with relative price (inequality (3)). The sensitivity of reach with respect to relative strengths and to relative prices diminishes as the brands are less similar (inequalities (4) and (5)).

Because we represent the intensity of each attribute for each brand as a real number in the unit interval, the coordinates representing the position of a brand in attributes space is always in the unit hypercube of dimensionality equal to the number of attributes. The maximum distance between any two points (corresponding to the diagonal of the hypercube) is the square root of its dimensionality — in this case the square root of the number of attributes. It is therefore natural to normalize the distances between brands' positions on the square root of the number of attributes. In this way, the model is not sensitive to the size of the chosen attribute set.

At the same time, we recognize that if two products are both very different from a third, how different they are from one another is not usually relevant to the consumers' brand choices. We therefore used a squashing function giving us a distance measure which made increases in small distances more important than the same increases in large distances. The function used in the model reported here was:

(6) $$\delta_{ij} = 2\tanh(|\Theta_i - \Theta_j|)$$

Because product differentiation need not have the same importance in all markets, we specify the differentiation effect as being determined by the distance between the products in attribute space and a *differentiation intensity parameter* (DIP) to be denoted as I_d. The differentiation effect expression is

(7) $$d_{ij} = e^{-(I_d \delta_{ij})^2}$$

where δ_{ij} is the distance between brands i and j in attribute space.

The effect of the relative strengths of two products is

(8) $$\Sigma_{ij} = \frac{\tanh\left(\left(\frac{\sigma_i}{\sigma_j} - 1\right) d_{ij} I_s\right) + \tanh(d_{ij} I_s)}{1 + \tanh(I_s)}$$

where I_s is the *strength intensity parameter* (SIP). The larger the value of the SIP, the higher the value of Σ_{ij} for any value of the strength ratio.

The effect of the relative prices is

(9) $$\Pi_{ij} = e^{-\frac{p_i}{p_j} d_{ij} I_p}$$

where I_p is the *price intensity parameter* (PIP).

The *reach function* is simply the product of the price effect and strength effect functions:

(10) $\rho_{ij} = \Sigma_{ij}\Pi_{ij}$

Preferences and the determination of market strength

The role of market strength in this model is partly, as we have seen, to determine the value of a brand's reach and also to determine the demand for each brand in each context. We turn now to the derivation of the measure of strength from the context-dependent attribute preferences.

We define a context-dependent attribute preference by an ideal value, a tolerance index and an importance index for each type of attribute included in the model. In The application to a market for spirits reported in section 5, for example, the four contexts are functional drinking, social, reward-seeking and novelty-seeking. The attributes of the brands sold in that market which were specified by the marketing professionals as uniqueness, specialness and expensiveness. Expensiveness is not the same as price or relative price since an "expensive" drink can sometimes be acquired (relatively) cheaply in a sales promotion.

A natural representation of ideal and tolerance is as a preference distribution function. This reflects the expectation that different consumers will have different ideals but that there is, at the same time, a central tendency to their ideal attribute values. For social purposes, for example, the tendency will be to value specialness quite highly and uniqueness very little. The dispersion of ideals for this purpose will be small and, in general, individual tolerance to deviation from the ideal is seen by the marketing professionals to be very limited. For this model, preferences are represented by a transform of the normal distribution such that the ideal value of an attribute is in effect the mean and the tolerance determines the variance. Importance is represented formally in a manner which takes advantage of the recognition that our preference-distribution function is not a probability function and is not required to integrate to unity.

In Figure 1, π_s on the vertical axis is the preference index determined by attribute value c for consumers in the context which entails undertaking activity s. The domain of π_s is the unit interval. The value of c is represented on the horizontal axis. Clearly, for a given ideal attribute value c^*, the dashed preference distribution entails more tolerance to deviations from the ideal than the solid-lined distribution. Moreover, the dashed distribution is more sensitive to attribute values below the ideal than to actual values above the ideal while the effect of deviations from the ideal is symmetrical for the solid-lined distribution

.

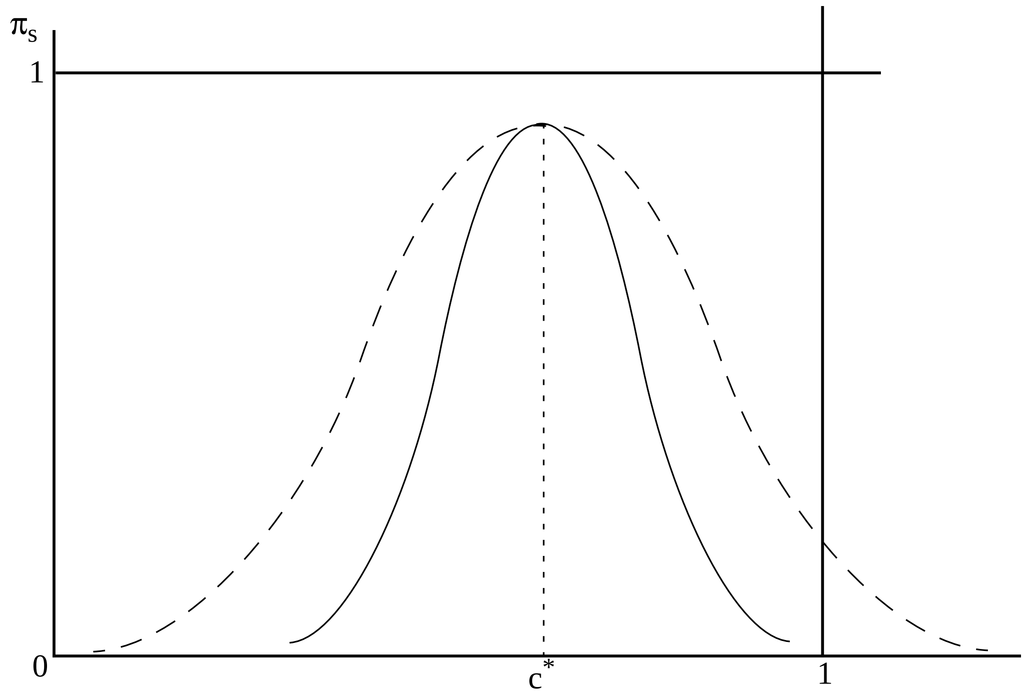

Figure 1: Preference distribution (same attribute ideal, different tolerances)

The preference index corresponding to the activity *s* context for brand *b,* denoted Γ_{sb}, is the product of the preference indices for actual attribute value associated with the brand. Formally,

$$(11) \qquad \Gamma_{sb} = \prod_{c \in C} \gamma_{sc}$$

where *C* is the set of defined attributes.

With this background, we turn now to the representation of importance.

It is easily seen that flatter and higher (in the sense of closer to 1) is the preference distribution, the smaller the effect if can have on the preference index of the brand for consumers undertaking the activity. If an attribute is completely unimportant, the preference distribution will be a horizontal line at the preference level equal to 1.

In Figure 2, we have the distributions differ only in their degrees of importance. Clearly, the flatter distribution is less important than the steeper distribution in that deviations from the ideal value entail preference indices closer to unity and, so, reduce the value of the context-dependent preference index for a brand by a lesser proportion

.

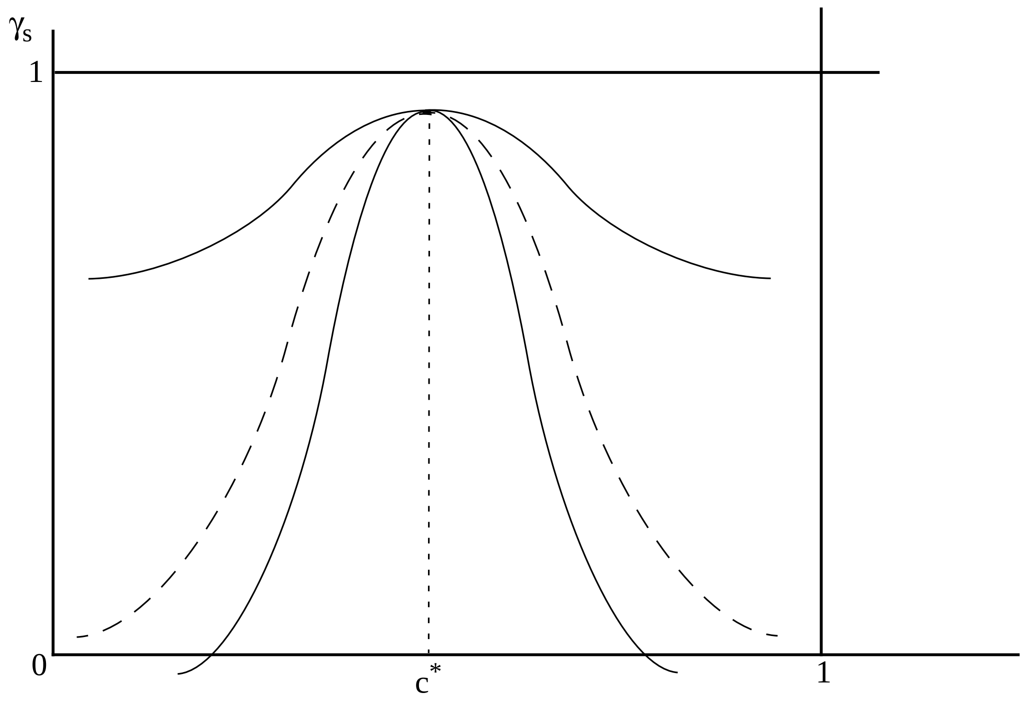

Figure 2: Preference distribution (same attribute ideal, same tolerances, different degrees of importance)

The preference distribution function used in the model reported here is a transform of the normal distribution since we are assuming that, at any time, individuals purchasing in a particular context are drawn at random from the population of potential purchasers and that there are random differences among them in respect of their ideal attribute values. In further developments, we will consider distributions of the importance and tolerance indices as well but, for the present, the preference distribution functional form is

(12) $$\gamma_{sc} = \frac{2m_c e^{-36\left(\frac{c^* - c}{t_c}\right)^2} + 1 - m_c}{1 + m_c}$$

where c is the value of attribute C, c^* is the ideal value, m_c is the index of the importance and t_c is the corresponding tolerance index. In practice, we have found that we get good empirical results with these models by mapping tolerance indices into the unit interval and importance indices into the [0, 0.5]-interval.

The derivation of the brand strength measure is now straightforward.

For strength of brand b in context s is

(13) $$\hat{\sigma}_{sb} = \frac{\Gamma_{sb}}{\sum_j \Gamma_{sj}}$$

which is the preference index for brand b expressed as a proportion of the sum of the preference indices for all brands in context s.

For any brand b, the market strength measure used in equation (8) is the sum of the context-related brand strengths as defined in equation (13) weighted by the proportion of sales accounted for by the corresponding contexts:

(14) $$\sigma_b = \sum_s \hat{\sigma}_{sb} w_s$$

We also use the context-related brand strength together with the values of the reach of brand *b vis a vis* all other brands to determine its notional demand index Δ_{sb} in context s:

(15) $$\Delta_{sb} = \hat{\sigma}_{sb} \prod_{i \neq b} \rho_{ib}$$

This value is not to be confused with an actual level of demand since the scale of these variable values is determined by the preference distribution parameters and not in any way by the data.

The total brand-demand index is the context-weighted sum of the values of the Δ_{sb}:

(16) $$\Delta_b = \sum_s \Delta_{sb} w_s$$

The market share is then determined by the brand-demand indices as if those indices were actual demands. That is,

(17) $$\mu_b = \frac{\Delta_b}{\sum_i \Delta_i}$$

4 Algorithms

If it is the case that the domain experts are dissatisfied with their present set of CDAP models, they may want some new models to adapt and work from. Typically these domain experts are fairly sure about the relevant properties in a particular market and the perceived characteristics of these properties for each product, what they are uncertain about is the number, identity and preferences of their customers. Thus there is a need for an algorithm which, given the relevant product characteristics, automatically searches for CDAP models that are consistent with the known data. An algorithm which we have found to be effective is described below - the Automatic CDAP Honing Engine (ACHE).

The heart of the procedure for determining a credible CDAP model from the sales data is a genetic programming algorithm. This is made more robust with a random search front-end to ensure a viable initial population of possible models and then a final hill-climbing evolutionary programming algorithm afterwards to tune the models found. This combination of three algorithms was found to produce qualitatively better solutions and work in a more robust manner than any of them separately. The robustness of this result is further enhanced by the implementation of the model in the logic-based programming language SDML (see [25]) which ensures that the individual algorithms and the integrated system are logically sound and consistent.

A binary-search hill-climbing algorithm optimizes the parameters of a single CDAP model, This algorithm can be used before and/or after ACHE at the discretion of the user.

4.1. Genetic programming module

Genetic programming (GP) differs from the familiar genetic algorithms in that the gene is a labelled tree rather than a string. The basic GP algorithm is:

1) Specify the possible branching and terminal nodes that the trees can be built from and the fitness function for evaluating them.
2) Generate an initial population of random trees of a given depth using these nodes.
3) Evaluate this population using the fitness function.
4) Find the best gene and, if it is good enough, stop.
5) Otherwise generate a new population of trees using one of two methods (according to a fixed proportion determined by the programmer):
   a) drawing pairs of trees randomly from the current population with a probability related to their fitness and producing two new offspring by choosing a random node in each and swapping the sub-trees that are rooted at these nodes (this is called tree-crossover) or,
   b) randomly choosing trees with fitness-related probabilities for propagation to the new population.
6) Go to step 3.

A basic account of GP with many applications is found in [14]. There are now many extensions and refinements of this technique.

In our case the tree-structure covered possible CDAP models. A gene was an instance of the following specification:

gene := *IP list*, *weight list*, *CDAP list*,

IP list := price intensity parameter, strength intensity parameter, differentiation parameter

weight list := list of non-negative numbers (of same length as list of CDAP states)

CDAP list = list of *CDAP specifications*, one for each CDAP state

CDAP specification := list of *preference specification*s, one for each property

preference specification := a triple of numbers: the ideal value, its importance and the tolerance to variation

The fitness function was the RMSE error of the predicted market shares compared to the actual shares over a sample period for the competitive set with a small discount to bias the algorithm in favour of models with fewer CDAP states.

Our crossover operator was constrained to produce only well-formed genes, i.e. if one chosen sub-tree was a preference specification the other would be also. Also if the domain expert had previously entered any trial CDAP models, these would be seeded into the initial population, so that variations of these would be tried along side the randomly generated ones.

4.2. Random Search

The range of attributes and attribute values that are at all acceptable to the candidate preference states can be small and CDAP models composed of such states can easily predict that no state demands any of the brands in the competitive set. This is a degenerate solution. The GP algorithm rapidly selects out such models, but if their proportion in the initial population is high then the effective variation in the initial population is restricted to the extent that the success of the GP module will depend critically upon the small subset of viable genes in the initial population.

In order to make the GP search more robust, a preliminary operation was devised which directly affects only the generation of the initial population. This operation randomly generates genes tests them for non-degeneracy until an initial population of the required size has been found which contains only viable genes. The test for viability is quicker than a full evaluation of fitness so this is relatively inexpensive in terms of computation time and results in a more comprehensive sifting of possible solutions by the GP algorithm.

4.3. Evolutionary programming

GP search algorithms are very effective in finding an acceptable solutions in large search spaces - i.e. they perform an effective satisficing global search. However, they are less effective in searching locally to determine the best variation of a good solution once it has been found. That is, while GPs are excellent satisficers, they are not good optimizers.

For this reason we added a last stage to the search algorithm which is essentially a multiple stochastic hill-climbing algorithm using an evolutionary programming technique. This works by keeping the fittest half of the population in each generation as well as generating near mutations of each of them. Thus each generation the original and the mutation is compared and the better ones selected.

We applied this to the population that resulted from GP algorithm using a mutation algorithm adapted to real-valued parameters which mutates parameters according to a normal distribution, so as to favour near mutations above far ones. The standard deviation of this distribution was ramped downwards throughout this phase so that large mutations would be tried before progressively finer ones, in a similar manner to a simulated annealing algorithm.

5 An Application

The IMIS system has been applied to a variety of markets in order to create CDAP models. One such application was to 28 monthly observations of scanner data covering supermarket sales of spirituous liquors in a large American conurbation. The data covered price and sales volumes by bottle of between 150 NS 200 different brands and bottle sizes. A United Distillers marketing expert decided upon the key attributes, measured on a Lickert scale, for each brand.

The IMIS system was given one focus brand from which it selected a competitive set of 4 other brands. The focus brand itself had a market small market share of less than five per cent on average in this competitive set although the brand itself has a significantly larger share worldwide. One issue for analysis is, therefore, why a generally successful brand is exceptionally unsuccessful in this one market.

The ACHE algorithm was set to work on the first 10 dates only in each of two experiments. In the first experiment, the algorithm settled on a model with a single preference state. In the second, we reduced the fitness discount associated with larger numbers of CDAP states and, as a result, identified four CDAP states. Each of these models was used to predict the future shares over the hold-out set. The model with a single CDAP state predicted market shares with an RMSE over all 5 brands of 7.1 per cent. The model with four CDAP states predicted the shares with an RMSE of 5.9 per cent. The graphs below exhibit the comparisons of the simulated and actual market shares of three of the brands. The other two brands accounted for less than 4% of the market and the models both ignored those brands by predicting zero market shares for them.

Increasing the number of CDAP states increases the accuracy of the system's simulation of market shares over whole data set. It also gives a very picture of the demand side of the market. This difference is discussed in section 5.3.

5.1. The single-CDAP-state model

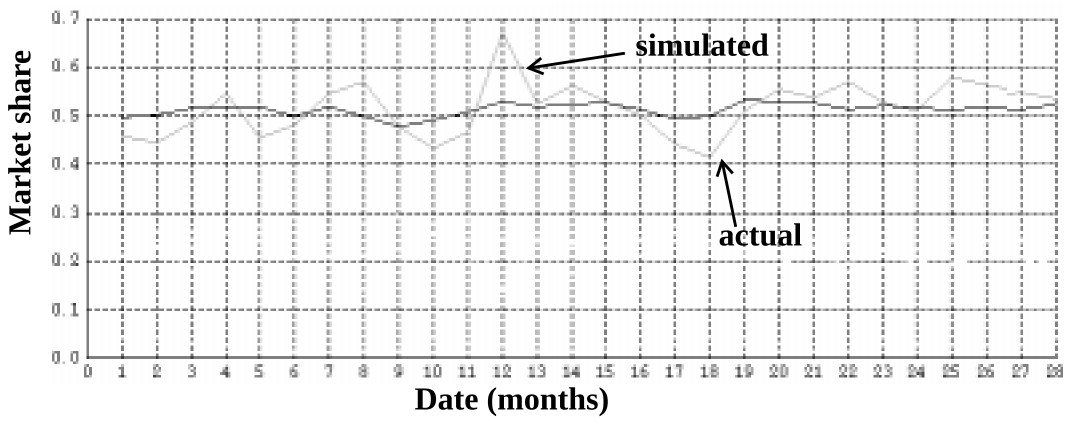

Figure 3: Real and Simulated Shares for Brand A (larger)

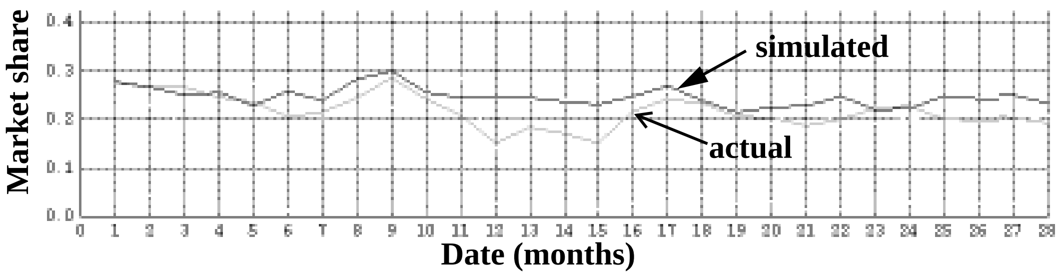

Figure 4: Real and Simulated Market Shares for Brand A (smaller)l

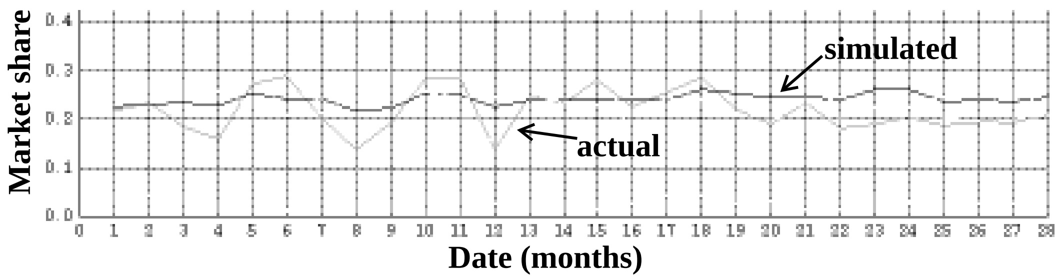

Figure 5: Real and Simulated Shares for Brand B (larger)

The characteristics of the CDAP state found in this experiment is summarised in Table 1. Expensiveness (the average price over the period) was most important brand attribute and the one to which consumers were least tolerant. This was followed by the amount of special ingredients and the size. They did not like imported spirits or special ones. They preferred low prices, but this was not critical and the consumers appeared tolerant of higher prices. In general terms, the single identified CDAP state want their chosen brands to be big, generally expensive and with lots of special ingredients.

Table 1: Qualitative Results from one-state model

Perceived Characteristic	Ideal Intensity of characteristic 0-1	Criticality (Sharpness of cutoff with deviation from ideal) 0-1	Intolerance to deviation from ideal 0-1
Relative Price	0	0.25	0.37
Expensiveness	0.87	0.62	0.62
Size	0.87	0.31	0.88
Importedness	0.12	0.44	0.62
Specialness	0.25	0.19	0.75
Fashionableness	0.38	0.44	0.75
Amount of Special Ingredients	0.87	0.56	0.75

5.2. The four-CDAP-state model

The simulated and real market shares generated by the four-state model are reported in the three figures 6 - 8. The preferences of each CDAP state are reported in Table 2. The CDAP-state accounting for the largest proportion of sales is state 1 which, as indicated in Table 2, is extremely concerned with the fashionableness of their purchase combined with some, rather more moderate, concern for specialness and distinctive ingredients. State 2 also wants its drink to be fashionable but is much less concerned with that characteristic and more concerned with cost and size (it wants the smallest sizes). In addition, State 2 has a moderately strong and well defined aversion to brands which appear to be imported. State 3 is much the same as State 2 in its preferences for cheap, small and rather special brands. The difference between States 2 and 3 is that, while both attach some importance to fashionableness, State 2 is intolerant of deviations from the most fashionable while State 3 is intolerant of deviations from the least fashionable. State 4, which is the second most important of the four CDAP states, is the least concerned with price but the most concerned with expensiveness. It is also just about equally concerned with fashionableness as State 1 but is intolerant of extremes in either direction. Size is much more important to State 4 than it is to State 1 and it wants the largest sizes.

In general terms, States 1 and 4, accounting for 70 per cent of the market between them, are up-market while States 2 and 3 are down-market.

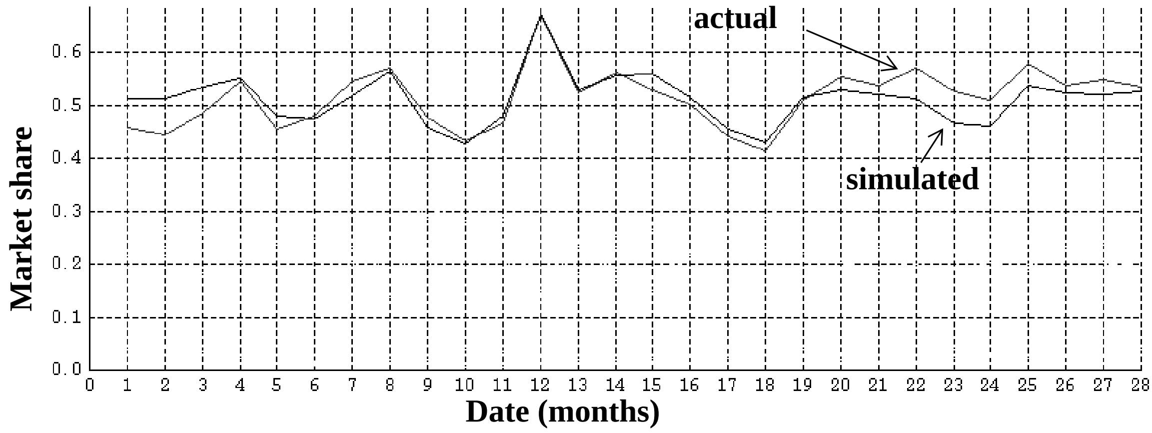

Figure 6: Real and Simulated Shares Brand A (larger)

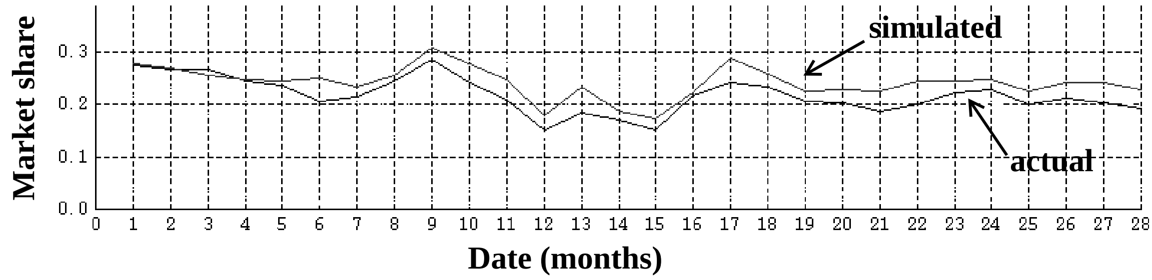

Figure 7: Real and Simulated Shares Brand A (smaller)

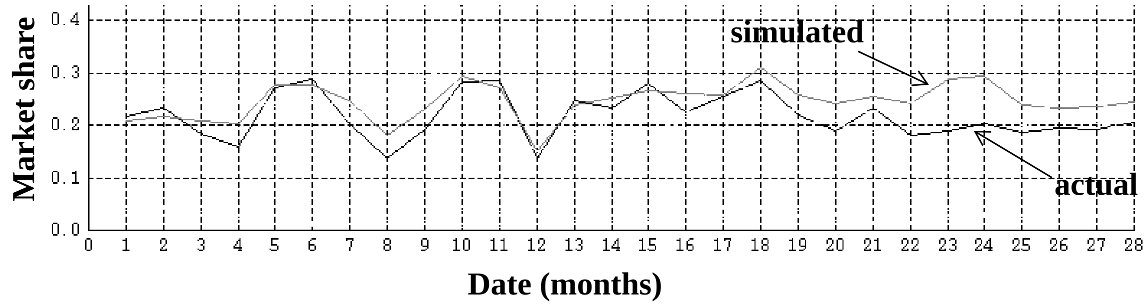

Figure 8: Real and Simulated Shares Brand B (larger)

Table 2: Qualitative results from four-state model
(criticality and intolerance indices in the unit interval in brackets)

Preference State Attributes	State 1	State 2	State 3	State 4
Notional Proportion of Total Sales	40%	15%	15%	30%
Relative Price	0.36 (0.8, 0.3)	0.4 (0.5, 0.8)	0.4 (0.6, 0.7)	0.4 (0.2, 0.7)
Expensiveness	0.36 (0.8, 0.3)	0.3 (0.6, 0.9)	0.3 (0.6, 0.9)	0.3 (0.8, 0.9)
Size	0.4 (0.2, 0.9)	0 (0.6, 0.8)	0 (0.6, 0.7)	9 (0.8, 0.5)
Importedness	0.4 (0.55, 0.4)	0 (0.6, 0.8)	0.4 (0.3, 0.5)	0.4 (0.6, 0.2)
Specialness	0.6 (0.7, 0.6)	0.7 (0.8, 0.7)	0.7 (0.8, 0.7)	0.4 (0.7, 0.6)
Fashionableness	1.0 (1.0, 0.8)	1.0 (0.4, 0.8)	0.1 (0.4, 0.8)	0.6 (1.0, 0.7)
Special Ingredients	0.4 (0.7, 0.7)	0.4 (0.9, 0.1)	0.4 (0.9, 0.1)	0.4 (0.7, 0.8)

5.3. A comparison of the two models

Two rather different marketing policies are indicated by the two models. The one-state model indicates that higher market shares will be achieved by making cutting the price of a usually expensive brand (though expensiveness is more important than relative price) while the four-state model indicates that relative price has both high criticality and intolerance for the 30 percent of the market accounted for by the two downmarket CDAP states. The one-state model indicates also that the market values moderately low fashionableness while the four-state model indicates that the largest CDAP state (in terms of the share of volume purchased) values extreme fashionableness very highly while the second largest state values moderate fashionableness and is less extreme in that valuation.

The four-state model also yields more plausible results with regard to specialness. In the one-state model, the special ingredients are ideally high and important while the specialness of the brand is ideally low and not very important (because criticality is low). In the four-state model, the ideal for special ingredients is the same for all CDAP states (0.4) but the importance is significant only for the up-market states. Specialness, on the other hand is important for all four CDAP states though one of the up-market states wants less specialness than do the others.

It is typical of applications of IMIS to such data sets that encouraging the definition of a larger number of CDAP states not only increases accuracy in tracking market shares but also, and more importantly, yields more detailed qualitative information which renders the results more plausible and easily interpreted.

In the application reported here, we can see how important the additional qualitative information is by comparing the two models' respective tables of preferences with the expert's specification of the brand characteristics in Table 3.

Table 3: Expert's specification of brand characteristics

characteristic	Focus (smaller)	Brand A (smaller)	Brand A (larger)	Brand B (larger)
expensiveness	1	0.2	0.9	0.0
size	0	0	1	1
importedness	0.9	0.5	0.5	0.1
specialness	0.7	0.3	0.3	0.5
fashionableness	0.3	0.5	0.5	0.3
special ingredients	0.3	0.9	0.9	0.3

Comparing the focus brand characteristics with the ideal characteristic mix of the one-state model, the focus brand would appear to miss the market by being seen as imported, the wrong size and lacking in special ingredients. On the other hand, it got fashionableness about right — though it is not all that far out from the other brands in that respect. In general, the focus brand's distinctive features were not well suited to the market described by a single CDAP state.

In the four-state model with preferences reported in Table 2, the focus brand's importedness looks offensive and important only to State 2 while both up-market CDAP states would be attracted by its size. Where the brand seems to lose the custom of the up-market CDAP states, accounted for 70 per cent of sales volume between them, is in its lack of fashionableness. The only CDAP state that would values its lack of fashionableness is State 3 which would be put off by the price and expensiveness.

6 Conclusion

The CDAP paradigm will support a variety of applications and can also rely upon and inform conventional market research techniques. We identify some of these extensions presently. It is also important to note, however, that CDAP models are likely to be inappropriate in some markets.

The markets in which CDAP analysis can make a substantive contribution are those where there can be a range of distinct purposes for making the purchase and where different attributes of the commodities purchased are required to fulfil the purposes of the purchase. In the case of alcoholic beverages, the range of purposes or needs which can be satisfied is large. But other branded consumer goods such as dishwashing detergents have only one purpose and are used only in one context: dishwashing. The application of CDAP functions to represent consumers in such markets would be interpreted as a distribution of household preferences and, as such, could be used to assess demands for different brands based on brand image alone. The distance measure and reach function would remain applicable in the obvious way.

While we do not assert that the CDAP paradigm is applicable to all markets, there are three natural developments which could be the subject of further research. One is the application to other types of market than markets for fast-moving consumer goods. Another is to determine the relevant attributes and to derive at least some indication of relevant CDAP states and corresponding functions by market research. The third is to use the CDAP paradigm for scenario analysis. These are considered in turn.

6.1. Examples of further applications

Two additional marketing issues that could usefully be modelled within the CDAP paradigm are industrial marketing and retail services.

In the case of industrial marketing, it is by no means uncommon for products made of different materials and with different engineering properties to satisfy the same user needs. A natural example is in materials. Steels and plastics have different properties relating to corrosion, heat resistance, expansion coefficients, weight, appearance and so on. Some of these attributes are clearly defined by the mathematics of physics and engineering while some such as appearance relate to customer perceptions. Marketing strategies could be geared to changing perceptions or making customers more aware of a largely unnoticed physical attribute or increasing the importance of one attribute for some purposes. CDAP models of a competitive set including some steels, other metals such as aluminium and some plastics could certainly be used to determine the consistency of the data and the suppliers' views of the reasons why customers and potential customers use the materials they do.

Retail services such as restaurants appeal to their customers very largely on the basis of image, atmosphere and similar attributes which amount to user perception rather than physical characteristics. There are, of course, as many different reasons for dining out as there are reasons for buying alcoholic beverages. As shopping increasingly becomes a leisure activity *per se*, then the image and atmosphere of other kinds of retail outlets will (or have) become similarly important. It is not hard to imagine relevant CDAP states for retail outlets. Functional (*i.e.*, shopping only to acquire specified goods) and recreational CDAP states come immediately to mind. There could well be different recreational or functional purposes that are relevant in different environments.

The usefulness of CDAP analysis in these markets is a matter for further research. We only note here that they are opposite ends of the spectrum from the dominance of customer perceptions to the dominance of physical characteristics.

6.2. CDAP specifications and market research

The attributes deemed to be relevant in the markets modelled in this paper were identified for us by collaborating marketing professionals with particular expertise of markets for alcoholic beverages. An alternative (though much more expensive) would be to determine the relevant attributes by market research. One obvious technique would be the long-established repertory grid (Kelly, 1955; Bannister, 1970; Bannister and Fransella, 1986; Stewart and Stewart, 1981). Moreover, to get some sense of the important CDAP states, conventional market research surveys could be extended to determine the activities being undertaken at the time a branded good was used or consumed.

6.3. CDAP simulation analysis

The constancy of the parameters of the CDAP functions has been taken for granted in the models reported here. Changing the relationships between prices and market shares is one of the objectives of marketing strategies. Such changes could result from affecting perceptions or affecting preferences. Changes in fashion are treated naturally as changes in the ideal value of an attribute while brand managers who follow fashion trends will seek to change the customers' perceptions of the brand to bring the perceived attributes in line with current fashion. Possibly, the importance of particular CDAPs or customers' tolerance to deviations of an attribute value from the ideal can be changed by marketing. These are issues that would have to be addressed by models which entail some representation of marketing strategies and their possible effects on customer perceptions and CDAPs. This, too, is a matter for further research.

Acknowledgements

The system and results reported here were obtained as part of the Intelligent Marketing Integrated System (IMIS) Project conducted in collaboration with United Distillers PLC as part of the Information Systems Integration Programme of the UK Engineering and Physical Sciences Research Council and the Department of Trade and Industry under contract number IED/4/8022. The programming language used in developing IMIS is SDML, designed and implemented by the Centre for Policy Modelling using the VisualWorks 2.5 development environment supplied by ParcPlace/Digitalk (UK) Ltd. The data was provided by United Distillers who have generously supported IMIS from its inception. Clive Sims of Competitive Analysis International suggested basing demand analysis on need-states which we formalized as CDAPs. We thank our colleague Steve Wallis and our collaborators in IMIS, Michael Campbell and Paul Ormerod for their encouragement and help in establishing the programme of which this paper is a result.

References

[1] Assael, H. (1992), *Consumer Behavior and Marketing Action*, (Boston: PWS-Kent Publishing Co).

[2] Belk, R.W. (1975), Situational Variables and Consumer Behavior. *Journal of Consumer Research* 2(December), 157-64.

[3] Campbell, M., S. Moss and C. Sims (1996), Rulebase-Driven Non-Linear Analysis of Competitive Structure. Centre for Policy Modelling *Technical Report 96-13*, Manchester Metropolitan University (http://www.fmb.mmu.ac.uk/cpm/cpmrep13.html).

[4] Collopy, F. and J. S. Armstrong (1992), Rule-based Forecasting. *Management Science* 38(10), 1394-1414.

[5] Crawford, J., A. Farquhar and B. Kuipers (1990), QPC: A Compiler from Physical Models into Qualitative Differential Equations. *Proc. National Conference on Artificial Intelligence* (Los Altos CA: AAAI).

[6] Deaton, A. and J. Muellbauer (1980), An Almost Ideal Demand System. *American Economic Review* 70, 312-26.

[7] Edmonds, B., S. Moss and S. Wallis (1996), Logic, Reasoning and A Programming Language for Simulating Economic and Business Processes with Artificially Intelligent Agents. In Ein-Dor, Phillip (ed.) *Artificial Intelligence in Economics and Management*. Boston, MA: Kluwer Academic, 221-230.

[8] Engel, J.F., R.D. Blackwell and P.W. Miniard (1995), *Consumer Behavior* (Fort Worth TX: The Dryden Press, 8th ed.).

[9] Forbus, K.D. (1984), Qualitative Process Theory. *Artificial Intelligence* 24(1), 85-168.

[10] Gordon, W. (1994), Meeting the Challenge of Retail Brands. *Admap* 29(3), 20-4.

[11] Hinkkanen, A., K.R. Lang and A.B. Whinston (1995), On the Usage of Qualitative Reasoning as an Approach Towards Enterprise Modelling. *Annals of Operations Research* 55, 101-137.

[12] Horsky, D. and P. Nelson (1992), New Brand Positioning and Pricing in an Oligopolistic Market" *Marketing Science* 11(2), 133-53.

[13] Huber, J., J.W. Payne and C. Puto (1982), Adding Asymmetrically Dominated Alternatives: Violations of Regularity and Similarity Hypotheses. *Journal of Consumer Research* 9(1), 90-98.

[14] Koza, J. R. (1993), Genetic Programming: On the Programming of Computers by Means of Natural Selection.. MIT Press, Cambridge:MA

[15] .Kuipers, B.(1986), Qualitative Simulation. Artificial Intelligence 29(3), 289-338.

[16] Lancaster, K.J. (1966), A New Approach to Consumer Theory. *Journal of Political Economy* 74, 132-157.

[17] Lilien, G.L, P. Kotler and K.S. Moorthy (1992), *Marketing Models* (Englewood Cliffs NJ: Prentice-Hall).

[18] Martilla, J.A. and J.C. James (1977), Importance-Performance Analysis. *Journal of Marketing* 41 (1), 77-79.

[19] Miller, K.E. and J. L. Ginter (1979), An Investigation of Situational Variation in Brand Choice Behaviour and Attitude. *Journal of Marketing Research* 16 (February), pp. 111-23.

[20] Morello, G. (1993), The Hidden Dimensions of Marketing. *Journal of the Market Research Society* 35(4), 293-313.

[21] Myers, J.H. and M.I. Alpert (1968), Determining Buying Attitudes: Meaning and Measurement. *Journal of Marketing* 32 (1), 13-20.

[22] Moss, S., M. Artis and P. Ormerod (1994), A Smart Macroeconomic Forecasting System. *Journal of Forecasting* 13 (3), 299-312.

[23] Moss, S. and B. Edmonds (1997), A Knowledge-based Model of Context-Dependent Attribute Preferences for Fast Moving Consumer Goods, *Omega*, 25, 155-169.

[24] Rust, R.T. and B. Cooil (1994), Reliability Measures for Qualitative Data — Theory and Implications. *Journal of Marketing* 31(1), 1-14.

[25] Wallis, S., B. Edmonds, B. and S. Moss (1995), The Implementation and Logic of a Strictly Declarative Modelling Language in A. Macintosh . and C. Cooper. (eds.), *Applications and Innovations in Expert Systems III* (Oxford: SGES Publications) 351-360.

[26] Wyatt G. R. Leitch and J. Scott, (1994), Qualitative Modelling in Economics. *Department of Economics Discussion Paper*, Heriot-Watt University.