

# Combining Evolutionary Computing Techniques to Find Credible Qualitative Descriptions of the Demand-Side of Markets

Bruce Edmonds, Scott Moss and Helen Gaylard,  
Centre for Policy Modelling, Manchester Metropolitan University,  
Aytoun Building, Aytoun Street, Manchester, M1 3GH, UK.  
Phone: +44 - 161 - 247 6482, Fax: +44 - 161 - 247 6802  
<http://www.fmb.mmu.ac.uk/cpm>

**ABSTRACT:** In many markets, customer preferences depend on the activity of the customer. The authors have previously demonstrated a model that allows the testing of the qualitative judgements of domain experts in spirits markets against relevant EPOS price and sales data. 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 a market for alcoholic beverages. The algorithm enables the automatic production of qualitative descriptions which are both consistent with observed data and deemed credible by the domain experts. The combination of these techniques is more robust and produces better results than any of these techniques separately.

## 1 INTRODUCTION - CONTEXT DEPENDENT ATTRIBUTE PREFERENCES

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 1997], however, have produced IMIS (Integrated Marketing Intelligent System) which embodies a computational model adopting a view from the marketing profession that, in some markets at least, preferences are more 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 same 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. Moss and Edmonds did not argue that CDAP models are in some sense better than statistical models - they have different purposes.

CDAP models support the incorporation of the qualitative judgements of domain experts in a way which enables us to test their consistency with the statistical record without measuring qualitative outputs directly. By “qualitative” we mean that some verbal expressiveness (e.g. conveying a notion of ordering within a class of well defined properties of objects) can be meaningfully mapped into a model. Such orderings can, but need not, be mapped into real number intervals. For each purchasing context, there will be some mapping between the perceived product attributes and its desirability. This distribution is approximated by a normal curve, parameterised in three ways: by the ideal intensity of the attribute; by the tolerance to deviation from this ideal, and the residual attractiveness when the ideal is not matched at all (i.e. the degree to which it is critically important). Consequently, for each purchasing context identified by the marketing professionals, we list the relevant attributes of the brands as well as the ordering of the intensity of such attributes.

Useful as that demonstration turned out to be, it was only applicable if the professionals knew the market well enough to posit useful characteristics of the demand-side. The professionals specified not only the brand attributes which they believed to be important in that market, but also the relevant contexts in which purchases were made. The first release of IMIS could modify but not fundamentally change the specifications of the professionals. The purpose of this paper is to report a technique to parameterise the CDAP functions automatically to find good descriptions of the demand-side.

## 2 A BRIEF DESCRIPTION OF THE UNDERLYING MODEL

For full details of the mathematical specification of the underlying model see [Moss and Edmonds 1997].

### 2.1 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. Given 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.

A focus brand is chosen. Then using algorithms driven by a knowledge-based system (as reported in [Campbell, Moss and Sims 1996]) we infer the chief price-competitors of this brand from data on prices and sales volumes for all brands in the data set. These algorithms have been developed and integrated in a somewhat *ad hoc* manner 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.

### 2.2 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 largely generated by taking from the shares of other (but possibly not all other) products. 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 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.

### 2.3 THE REACH FUNCTION

In order to capture these ideas in a model we define a function called *reach*. This 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 when there is less similarity between the products. The market strength is estimated by the extent to which the perceived characteristics of the brands meet the preferences of the groups of customers in the respective purchasing contexts.

## 3 THE 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 properties that are relevant to a particular market and the perceived characteristics of these properties for each product. What they tend to be uncertain about is the number, identity and preferences of their customers. Thus there is a need is for an algorithm which given the relevant product characteristics automatically searches for CDAP models that are consistent with the know data for the given competitive set. An algorithm which we have found to be effective is described below - the Automatic CDAP Honing Engine (ACHE).

The heart of the algorithm is a genetic programming (GP) algorithm. This is made more robust with a random search front-end to ensure a viable initial population and a hill-climbing evolutionary programming algorithm afterwards to tune the resulting models. This combination of three algorithms was found to produce qualitatively better solutions and to work in a more robust manner than any of them separately, though we have not space to exhibit the comparisons here.

### 3.1 GENETIC PROGRAMMING MODULE

Genetic Programming is a technique similar to the genetic algorithms, however the gene is not a string but a labelled tree. The basic technique involves:

- 1) specifying the possible branching and terminal nodes that the trees can be built from and a fitness function for evaluating them,
- 2) generating an initial population of random trees of a given depth using these nodes,
- 3) evaluating this population using the fitness function,
- 4) Find the best gene and see if it is good enough, if so stop,
- 5) generating a new population of trees using one of two methods (according to a fixed proportion determined by the programmer) - *firstly*, 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) and, *secondly*, merely randomly choosing with a probability related to its fitness and propagating it to the new population.
- 6) repeat steps 3 and 4 until either an acceptable gene is found or a set limit is reached.

A basic account of this technique with many applications is found in [Koza 1992]. There are now many extensions and refinements of this technique. In our case the tree-structure covers possible CDAP models. Thus a *gene* is an instance of the following specification. The three parameters are needed to help scale the qualitative input.

*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 CDAPs)

*CDAP list* = list of *CDAP specifications*, one for each preference group

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

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

The fitness function is 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 groups. Our crossover operator was constrained to only produce 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 will be mixed with the randomly generated ones.

### 3.2 THE RANDOM SEARCH FRONT-END

Typically for many markets the range of characteristics that are at all acceptable to the candidate preference groups can be small. CDAP models composed of such groups can easily predict that none of the customers have any demand for the products at all. This is a degenerate solution. The GP algorithm rapidly selects out such models, but if their proportion in the initial population is high then this greatly restricts the effective variation in the initial population and hence the success of the GP module can critically depend upon the particular mix of characteristics found in this small subset of viable genes in the initial population. In order to make the GP search more robust, a random search was inserted before it. This affects only the generation of the initial population. The algorithm randomly generates genes (of a correct structure) and tests them, until an initial population of the required size of viable genes has been found. The test for viability is quicker than a full evaluation of fitness so this is relatively inexpensive in terms of computation time in return for a more comprehensive shuffling of possible solutions by the GP module.

### 3.3 THE EVOLUTIONARY PROGRAMMING OPTIMISATION BACK-END

GP search algorithms are quite robust, in that they are good at finding an acceptable solution in large search spaces - i.e. they perform an effective satisficing global search. They are less good at narrowing down to the best variation of a good solution once they have found it, i.e. they are not so good at optimising. 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 [Bremermann et al. 1966]. 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.

#### 4 AN APPLICATION TO A MARKET FOR SPIRITS

As an example of the output of this algorithm we applied the technique to monthly EPOS data from a large market for spirituous liquor in the US, with about 80 different brands, coming in two sizes. The data covered 160 lines over 28 monthly observations. A marketing expert decided upon the key dimensions for the market and the characteristics along these dimensions for each line.

The IMIS system was given one focus brand from which it selected a competitive set of 4 other lines. The ACHE algorithm was set to work on the first 10 dates only. The algorithm settled on a single preference group. This group was then used to predict the future shares which it did with an overall RMSE over all 5 lines over all 28 dates of under 4%. Graphs of three of the brands are shown below in figure 1 and figure 2 (the fourth and fifth accounted for less than 4% of the market for which the model gave zero forecasts). This run was better than the average performance of the system, but this accuracy shown is far from exceptional for it.

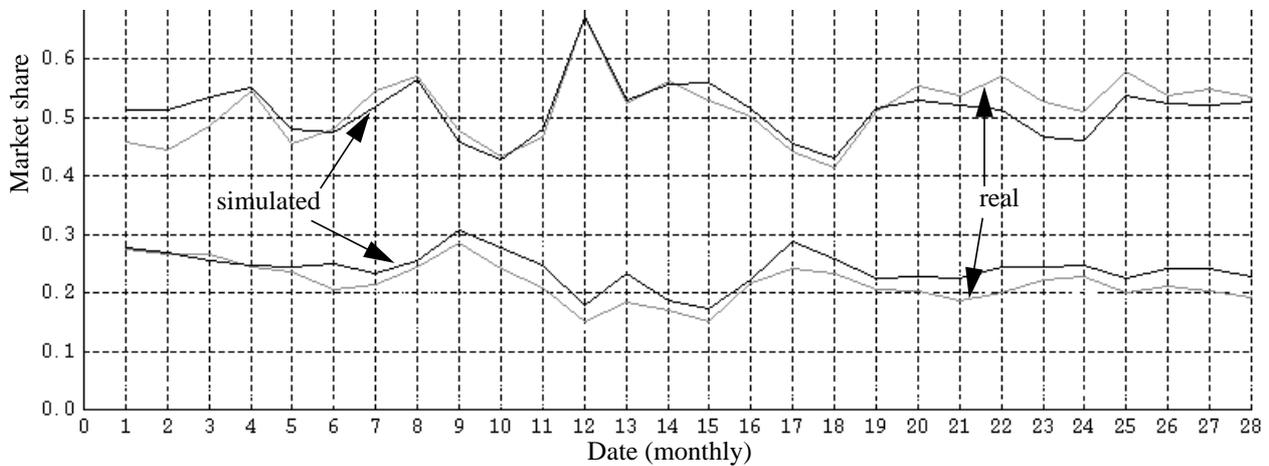


Figure 1: Real and Simulated Shares Brand A and B

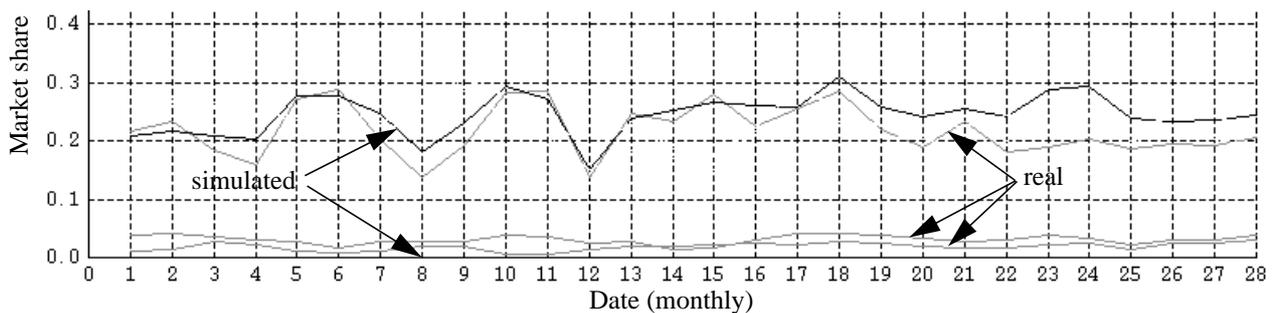


Figure 2: Real and Simulated Market Shares for Brands C, D and E (simulations for D and E are flat and zero)

In [Moss and Edmonds 1997], we compare the performance of the knowledge based IMIS model with linear regressions. The IMIS model generally has a smaller error than the linear regression over the when parameterised over only few dates. Typically if the linear regression is parameterised over more than half the data it does better. However output from this is not easily interpretable by marketing professionals, and number of dates required to acceptably parameterise it means that the technique can not be used on shorter subsections of the data to help identify what might have changed in the perceptions of the customers.

More interesting are the characteristics of the groups it found. The algorithm found four preference groups. This is summarised in table 1, below. Here only the ideals which were significant are shown (in terms of discovered tolerance and importance), the ones in bold are particularly significant for the group. The ideals are given from 0-10 which is scaled to span the maximum range of characteristics of the brands in the given competitive set.

Preference Group Attributes	Group 1	Group 2	Group 3	Group 4
Notional Proportion of Total	40%	15%	15%	30%
Relative Price	4	4	4	
Expensiveness	0	<b>3</b>	<b>3</b>	<b>3</b>
Size	4	<b>0</b>	0	
Importedness				
Specialness		7	7	
Fashionableness	<b>10</b>	10	1	<b>6</b>
Special Ingredients	4			<b>4</b>

Table 1: The significant ideals of the discovered preference groups

One can see that it came up with sensible results. All the groups preferred lowish relative prices, inexpensive brands, and smaller size. The groups could be distinguished as follows: group 1 - fashionable & cheap; group 2 - fashionable and special; group 3 - unfashionable, and special; and group 4 -fairly fashionable without many special ingredients. A marketing professional might label these are cheap & fun, special, fuddy-duddy and regular.

## 5 CONCLUSION

The CDAP paradigm will support a variety of applications and can also rely upon and inform conventional market research techniques. 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. We have shown a robust technique for searching for plausible specifications of such a demand-side, given sufficient EPOS data, and the relevant properties of the brands concerned. This enables the automatic determination of plausible models of the demand-side of markets for FMCGs, and does so in a manner that is readily interpretable in the qualitative language of marketing professionals.

## 6 REFERENCES

- Bremermann, H.J.; Rogson, J.; Salaff, S.; 1966. Global properties of evolution processes. In Pattee, H.H. (ed.). *Natural Automata and Useful Simulations*. Spartan Books, Washington, pp. 3-42.
- Campbell, M.; S. Moss; C. Sims 1996. Rulebase-Driven Non-Linear Analysis of Competitive Structure. Centre for Policy Modelling Technical Report 96-13. Accessible at <http://www.fmb.mmu.ac.uk/cpm/cpmrep13.html>
- Koza, J. R. 1993. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge/MA.
- Moss, S.; M. Artis; P. Ormerod 1994. A Smart Macroeconomic Forecasting System. *Journal of Forecasting* Vol. 13, pp. 299-312.
- Moss, S.; Edmonds, B. 1997. A Knowledge-based Model of Context-Dependent Attribute Preferences for Fast Moving Consumer Goods, *Omega* Vol. 25.