20 February, 1995 (saved at 1046)

A Formal Preference-State Model with Qualitative Market Judgements

SCOTT MOSS AND BRUCE EDMONDS
CENTRE FOR POLICY MODELLING
MANCHESTER METROPOLITAN UNIVERSITY

Formal models of market demands derived from the economic theory of choice have four deficiencies for the modelling of some FMCG markets.

1) For many FMCGs — especially food and drink products — customer preferences depend on the purposes for which purchases are made and each individual can have different purposes in mind when purchasing at different times. A theory of household demand (with one preference function for each agent) is therefore an inappropriate construct for the analysis of these markets.
2) The theory has it that, income effects aside, cross-price elasticities of demand are symmetric as between pairs of products. In practice, asymmetries are present and important.
3) Although brand data in these markets is typically much better than customer data, the theory implies a great deal about the customers and these implications dominate the predictions and forecasts generated by the models predicated on the theory.
4) Marketing experts typically have a qualitative understanding of important aspects of their markets and these cannot be represented by means of utility functions or partial orderings based thereupon.

In this paper, we set out an alternative modelling procedure for markets in which demand asymmetries are possible, qualitative aspects of demand can dominate, preference-states determine the preferences of individuals and relatively poor customer data is not a bar to modelling demand when there is good brand and product data. We demonstrate that models entailing these conditions can be used to explain observed qualitative characteristics of products, markets and sources of demands as well as good, numerical data such as EPOS data. Indeed, we will argue that models that rely on qualitative judgements and closely track actual numerical data series are more useful than conventional models which rely on numerical data to track numerical data.

1 The deficiencies of the economic theory of choice

Demand asymmetry.

Conventional choice theory has it that only price ratios are important in pairwise comparisons of the demands for goods. The level of each price is unimportant except in relation to other prices and incomes. All demand models must be homogeneous of degree zero. Ignoring income effects, the effect of changing any price ratio is the same whether it is the result of a change in the numerator, the denominator or some combination of both.

Consider three products: a litre bottle of some modest whisky, a 70cl bottle of the same whisky and a 70cl bottle of some rather more special whisky. It is at least plausible that a reduction in the price of the litre bottle will induce consumers to substitute litre for 70cl bottles of the same whisky while an increase in the price of the 70cl bottle will induce consumers to substitute a slightly more special whisky and not move up to the litre bottle of the same whisky.

The source of this asymmetry is not related particularly to the volume of whisky per penny. Reducing the price of a mellower but more expensive whisky will induce a substitution from a rougher, cheaper whisky. But increasing the price of the rougher whisky will induce substitution to other, still cheap brands of whisky. This is not an asymmetry brought about by income effects. It is a straightforward asymmetry in the substitution effect among the attributes or characteristics of the alternative products. In the first example, the relevant attributes were the “good value” of the litre bottle and the “specialness” of the competing brand. These determined their market strengths relative to the 70cl bottle of the same brand as the litre bottle. In the second example, the relevant attributes were the mellowness and the expensiveness of the alternative brands. In general, asymmetries seem likely to occur where there are three products and two or more product attributes.[1]

State-based preferences.

It has been well known for at least twenty years that the purchasing situation affects consumers’ preferences.[2] Recently, the suggestion that consumers have different preferences in buying for different purposes has gained some currency among practising marketing managers.[3]

Qualitative judgements.

In the growing literature on qualitative modelling, the qualitative nature of the models is generally concerned with the relative values of variables. That is, the values of a qualita-

---

[1]Asymmetries seem less of a problem in the marketing literature than in the economics literature. Compare, for example, Horsky and Nelson (1992) with the seminal article by Deaton and Muellbauer (1980).

[2]*cf.* Belk (1975), Miller and Ginter (1979), Srivastava, Alpert and Shocker (1985)).

[3] See Morello (1993) and Gordon (1994).

tive variable would be pretty much restricted to “up”, “down”, “larger” or “smaller”. Marketing professionals use a richer set of qualitative values to describe the attributes of products. Two that are important in the market modelled below are “uniqueness” and “specialness”. Some products will be perceived as being more or less special or unique than others.

In practice, we represent the degree of specialness or uniqueness as numerical values which have no intrinsic meaning. But “more special” or “less special” are simply “larger” and “smaller” applied to the numerical indices of the qualitative variables. Provided that we are interested in a partial ordering with some sense of the degree of difference among the qualitatively-valued variables, then the results obtained for the qualitative models of, for example, Leitch and Wyatt (1994) apply here.

Uneven data reliability.

We touch here on a methodological point to be explored in greater depth below.

Qualitative data is, in its nature, imprecise though not necessarily unreliable. The difference between the accuracy and the precision of numerical data is well known. There is no problem in collecting and calculating statistics of great precision and limited accuracy. Where data about whole economies is concerned, a major problem is that successive revisions of data can change the shape of our appreciation of the economic circumstances of the past out of all recognition. Spectacular examples include the falling UK savings ratio of the early 1970s (it turned out to have risen) and the disastrous trade figures on the eve of the 1970 British election which are widely credited with the defeat of the Wilson government and (the balance of trade improved).

The problem with qualitative data is the difficulty of comparing simulated results with actual results. It is therefore important that our simulation results include some outputs corresponding to numerical data series as well. It is clearly preferable to try to track the most reliable series where by reliable we mean the series which are accurate to the greatest degree of precision. As Wyatt, *et. al.* have shown, less precision can sometimes improve accuracy. However, the model reported below is tested by attempting to track weekly EPOS data from UK supermarkets. This date is extremely reliable. We use it to reach qualitative judgements about the characteristics of the demand side of the market for which we have no reliable data.

We do not claim precision but we do claim plausibility in the characterization of the demand side. The justification of plausibility is the assessment of the results by marketing professionals with expertise of the market we modelled.

2 An alternative modelling paradigm

Our approach in general is to represent products by their characteristics, the intensities of the respective characteristics, their market strengths and their prices. These three elements of the products determine their market shares. We do not, at this stage, address the issue of absolute demand determination.

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. 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 characteristics. 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.

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 state 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 characteristic 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 characteristics vectors is $|\Theta_j - \Theta_i|$.

Formally,

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

where

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

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

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

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

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 characteristic for each brand as a real number in the unit interval, the coordinates representing the position of a brand in characteristics space is always in the unit hypercube of dimensionality equal to the number of characteristics. The maximum distance between any two points (corresponding to the diagonal is the square root of its dimensionality — in this case the square root of the number of characteristics. It is therefore natural to normalize the distances between brands' positions on the square root of the number of characteristics. In this way, the model is not sensitive to the size of the chosen characteristics 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) $$\delta_{ij} = 2\tanh\left(\left|\Theta_i - \Theta_j\right|\right)$$

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 characteristics 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 characteristics space.

The effect of the relative strengths of two products is

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

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 characteristics of a model of consumer demand which reflects current thinking of marketing professionals and which meets the most stringent criteria of rigour and methodological rectitude will have the following characteristics:

- The demand side will be represented by preference-states rather than households or demographic groups

- Qualitative output will describe the demand side of the market for which there is little or no hard evidence.

- Some of the output must be numerical so that it can be compared with reliable numerical data series.

We define a preference state by an ideal value, a tolerance index and an importance index for each type of characteristic included in the model. In the next section, for example, we report an application to a market for spirits with four preference-states: functional, social, reward-seeking and novelty-seeking. The characteristics of the brands sold in that market which were specified by the marketing professionals. They were uniqueness, specialness and expensiveness[4].

A natural representation of ideal and tolerance is as a preference distribution function. For this model, we developed a transform of the normal distribution such that the ideal value of a characteristic 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 of the preference-state *s* determined by characteristic value *c*. The domain of π_s is the unit interval. The value of *c* is represented on the horizontal axis. Clearly, for a given ideal characteristic 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 characteristic 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.

The preference index of preference-state *s* for brand *b*, denoted Γ_{sb}, is the product of the preference indices for actual characteristic value associated with the brand. Formally,

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

where *C* is the set of defined characteristics. In the model reported here, $C = \{$*uniqueness*, *specialness*, *expensiveness*$\}$.

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

---

[4]Expensiveness is not the same as price or relative price since an “expensive” drink can sometimes be acquired (relatively) cheaply in a sales promotion.

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 the preference-state. If a characteristic 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 preference-state preference index for a brand by a lesser proportion.

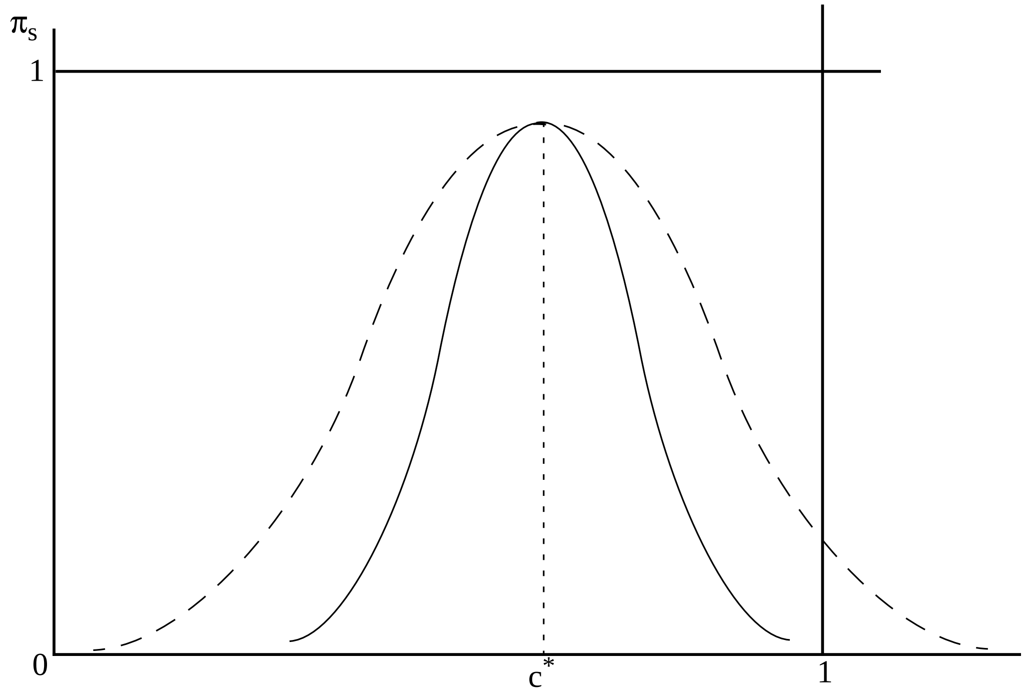

Figure 1: Preference distribution (same characteristic ideal, different tolerances)

The preference distribution function used in the model reported here is

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

where c is the value of characteristic 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.

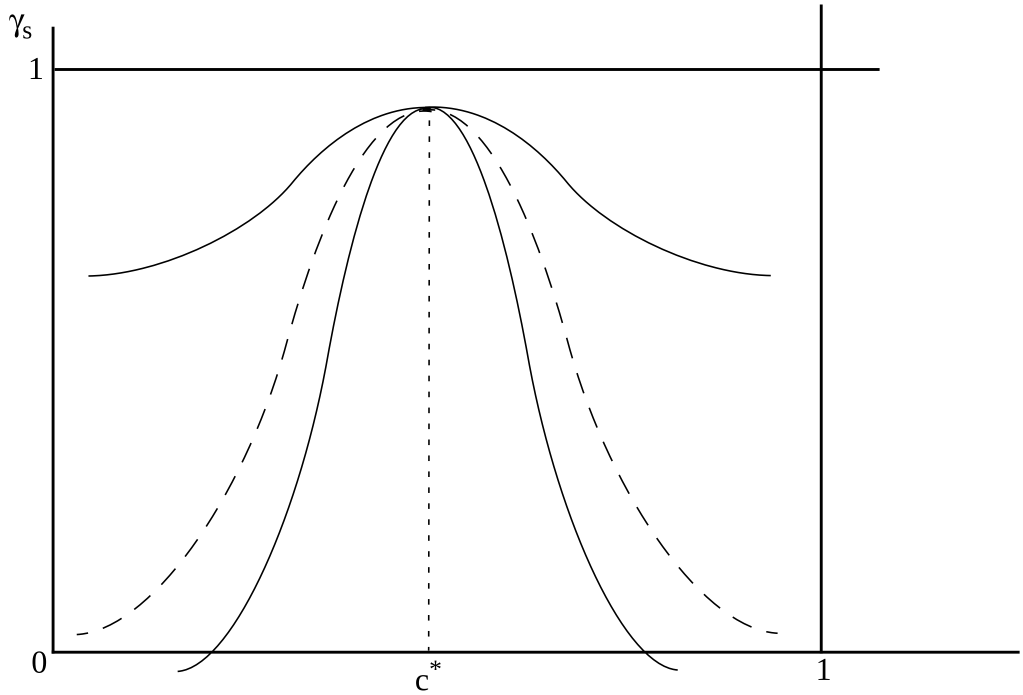

Figure 2: Preference distribution (same characteristic ideal, same tolerances, different degrees of importance)

The model is solved for volume shares in the manner indicated in Figure 3. The basic elements of the equation are reproduced with their equation numbers as are five additional expressions — (S-1) to (S-5) — defining appropriate combinations of the elementary preference and reach variables. Expression (S-1) defines the relative desirability of each brand for preference-state *s* as the value of the preference function for the brand divided by the sum of the preference values for all brands for that preference-state. The relative desirabilities have two uses. One is in the calculation, in expression (S-2), of the market strength of each product. This is the average of the relative notional preference-state demands weighted by the proportions of total sales accounted for by each preference-state. The second use of the relative desirabilities is in expression (S-3) where they are multiplied for each preference-state by the reaches of all products ($\rho_{ii} \equiv 1$) to yield the levels of the notional demands in each preference-state for each brand. Summing over preference-states in expression (S-4) fives us an index of total demands for each brand. It is then a simple matter in expression (S-5) to calculate the market share of each brand as its own demand index divided by the sum of all demand indices.

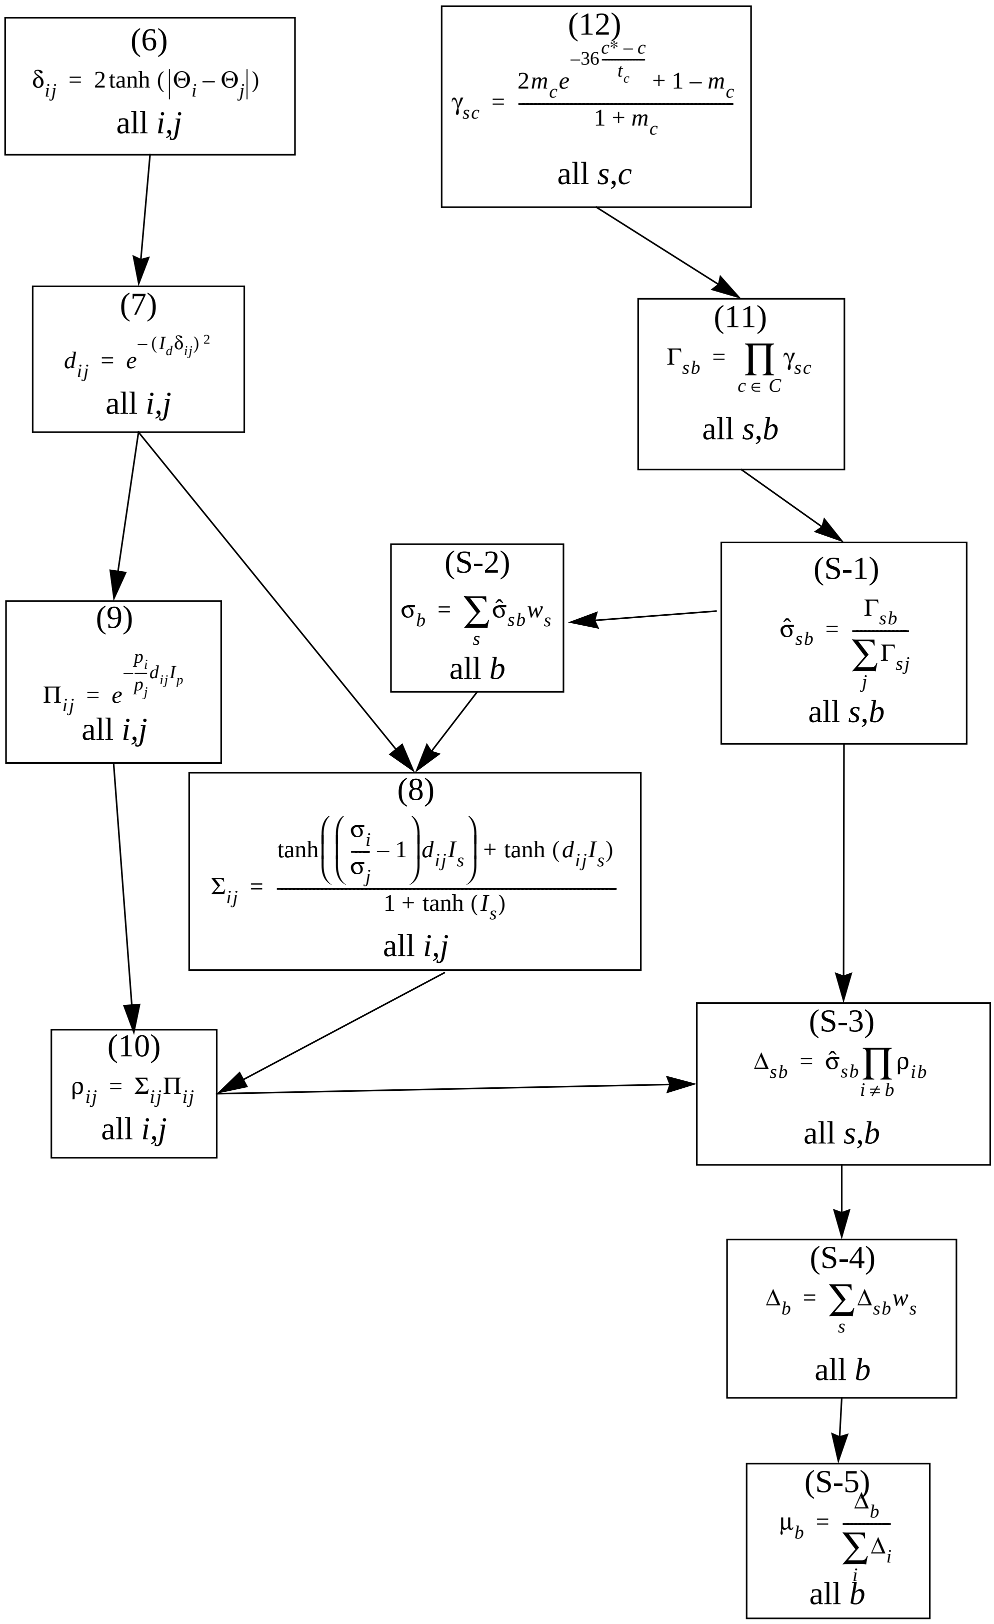

Figure 3: Solution flow-chart

3 Parameterizing the model

The algorithm described in the preceding section will give yield a unique solution for every set of meaningful (in this case, real, positive) parameters. The next issue is the choice of parameter values.

Our objective is to determine the parameter values that will minimize some measure of the difference between simulated shares calculated from actual prices and actual market shares. A close correspondence between the tracks of simulated and actual shares over the data period will give confidence in the validity of the model as the basis for counterfactual analyses of the market and market strategies.

Specifications of preference-states and estimates of the parameters of the corresponding preference functions are derived from qualitative judgements elicited from the marketing professionals. An example of this procedure is described in some detail in section 4.

Though the direction of changes in the values of the three intensity parameters (PIP, SIP and DIP) have obvious qualitative meaning, the levels of their values do not correspond to any meaningful empirical referent. We do know from numerical analysis that the RMSE with respect to market shares has multiple minima corresponding to different values of the DIP though the values of those minima vary with the PIP and SIP. High DIP values imply that the brands have small effect on one another. Consequently, high DIP values make it relatively easy to find other parameter values which get the shares about right at any one date but then changes in relative prices will have small effects on shares. Conversely, low DIP values make share values depend more critically on other parameters — and so make it difficult to get accurate share simulations — but easier to get appropriate amplitudes in share variations over time.

The algorithm used to parameterize the model reported below applies binary search to the three intensity parameters in turn and then to the proportions of total demands accounted for by each of the defined preference-states. This is followed by allowing for a single change in an ideal characteristic value followed by changes in certain importance and tolerance parameters — though in the results reported here the algorithm never got that far before finding an acceptable solution. The objective of the algorithm was to minimize the RMSE of all, unweighted market shares.[5].

The application of binary search is standard. Initial values for each parameter and an initial percentage change are specified by the modeller. The effect of the first change is assessed. If it results in increased accuracy, a further change is made in the same direction but the percentage change is halved. If the effect is diminished accuracy, then the same percentage change is made relative to the initial value but in the opposite direction. Again, an improvement implies a further change by half the amount in the same direction. If the initial percentage change in both directions results in reduced accuracy, then halve the percentage change from the initial value and try that on either side. Whenever accuracy is improved, take the improved value as the base and continue to search on either side of it in the same manner to find further improvements. The percentage by which the value is changed is not increased after each halving, As a result, there is a limit to the change in any one parameter before every other parameter is changed.

Characteristic ideals were changed for the preference-state accounting for the largest demand of the brand with the largest error in its simulated share. If the error is positive (simulated exceeds actual share) and the preference-state also demands another brand with a different value of at least one characteristic, then change the ideal in the direction of the actual value of the same characteristic of the other brand. If the error is negative, then move the ideals towards the brand with the largest error. Because the ideal characteristic values are most clearly identified by informants, these are changed only by small amounts and only once in each cycle. Moreover, they are only changed when the ideal value is between the two actual values and are never changed so that the ideal is outside that range.

The tolerance and importance indices of the characteristics are changed more freely since their values are not seen with the same conviction by informants. There are four cases. Consider first positive errors so that a reduction in demand is wanted. If there is an ordering of the actual characteristic of the largest-error brand is closer to the ideal of a preference-state than the brand with the largest demand by that preference-state (ignoring the brand with the largest error). then increase the importance of the characteristic and reduce the preference-state's tolerance to deviations from the ideal. The converse case follows naturally. A bisection algorithm is then applied to these changes until the best set of tolerance and importance indices, given the other parameter values, is found.

In all cases, the algorithm moves from one parameter to the next whenever the change in the absolute value of the largest error is less than some minimum precision — in the

---

[5] Other measures such as MAPE or the use of share-weighted error measures would have been possible. While further elaboration of the solution algorithms along these and other lines will doubtless yield improvements, our purpose here is to demonstrate that the overall methodology is sound and useful.

model reported here 0.005. As already indicated, the algorithm halts when the error measure is less than the minimum precision.

4 An application

In this section we describe a model that was constructed to represent a market in which United Distillers (UD) is active. The EPOS data for the brands included in the model were provided by UD. The qualitative judgements about the demand side were provided by Clive Sims, then UD's head of technical marketing.

Initial specification of parameter values

As previously indicated, the demand side comprises four “preference-states”: the functional, the social, the self-rewarding and the novelty-seeking states. Each preference-state is defined by the ideal values of product characteristics, the importance of each such characteristic and the tolerance to deviations from the ideal. The ideal characteristics for each preference-state as described by Sims are reported in Table 1. The consumer in the functional preference-state is indifferent to the characteristics “uniqueness” and “specialness” but wants the purchase to be as inexpensive as possible. When making purchases for social consumption, the consumer wants a drink that will have wide appeal and, so, will not be in any way unique, but will be special and moderately expensive. The consumer who is making a purchase in search of novelty will want maximal uniqueness, moderate specialness and is indifferent to expensiveness. Finally, reward-seeking consumption is best satisfied by rather special, rather expensive drink but the uniqueness for this purpose is irrelevant. These judgements were mapped into the unit interval as given by the numbers in Table 1.

Table 1: Desired characteristic intensities for each purchase type

Purchase type (% of sales volume)	Uniqueness	Specialness	Expensiveness
functional (40%)	?	?	0
social (40%)	0	1	0.4
novelty-seeking (5%)	1	0.6	?
reward seeking (15%)	?	0.8	0.8

Table 2: Preference-State Specifications

		Functional	Social	Reward	Novelty
Uniqueness	Ideal	0.5	0	0.2	1.0
	Importance	0.2	0.8	0.2	1.0
	Tolerance	1.9	1.2	2.0	0.4
Specialness	Ideal	0.5	1.0	0.8	0.6
	Importance	0.2	1.0	1,0	0.6
	Tolerance	1.9	1.2	0.4	1.2
Expensiveness	Ideal	0	0.2	0.8	0.5
	Importance	1.0	0.8	1.0	0.2
	Tolerance	1.4	2.0	1.2	2.0
Price	Ideal	0	0	0	0
	Importance	0.8	0.8	0.4	0.4
	Tolerance	0.2	0.6	1.0	0.8
\percentage of demand	Calculated (specified)	27.5 40	27.7 40	24.9 15	19.9 5

The expert's estimates of the degrees of importance of the different characteristics in relation to the different preference-states and the level of tolerance to deviations from characteristic ideals are given in Table 2. In that table, these values are specified in the unit interval with indifference over the ideal value of a characteristic being represented as 0.2. The reason we did not set the importance indices for indifference to zero will be explained presently. In the model, importance indices can take any positive real value. Expressing them here in the unit interval simply gives some sense of relative importance.

The marketing professionals suggested that functional purchases account for about 40% by volume (numbers of bottles) of this market, that social purchases account for perhaps another 40 percent, self-reward accounts for some 15 per cent of purchases with the remaining 5 per cent being purchases by consumers seeking novelty. In setting these as the initial values for the preference-state proportions we found the model did not solve easily or quickly. In setting all of the initial proportions to one-quarter, however, the model was solved easily with the preference-state proportions in the order suggested by the marketing professionals but with the proportions much closer together. Social and functional purchases accounted for 27-28 percent, reward-seeking about 25 percent and novelty-seeking just under 20 percent (see Table 2).

Estimated intensity parameter values

The levels of the three intensity parameters have no natural meaning and no meaning at all to the marketing professionals. Changes or differences in the values do have meaning. An higher value of the differentiation intensity parameter (DIP) implies less competition among the brands in the market. A higher value of the price intensity parameter (PIP) indicates that relative prices are more important in the determination of brand shares. A higher value of the strength intensity parameter indicates a greater importance for characteristic preferences in share determination.

The initial values of these parameters were necessarily arbitrary though they did reflect our understanding of functional form of the relations constituting the model and how they were intended to interact. The parameter values used (with their initial values in parentheses) were DIP = 0.125 (2.0); PIP = 2.11 (1.5), SIP = 0.74 (1.5). In fact, the PIP and SIP did not change the RMSE by more than minimum precision and, so, were only changed once each. The DIP value, however, was more effective in reducing the RMSE. It is plausible that the value of the DIP generated by the solution algorithm indicates that the market is highly competitive. To reach this conclusion in any definitive or formal way on the basis of preference-state models such as ours will clearly require experience in applying the modelling framework to a wide range of individual markets.

5 Simulation results for the model

The simulated shares reported in this section were calculated from the parameterized model and the actual price data over the 96-week data period. The model was parameterized by searching for the set of parameters yielding an unweighted, one-period RMSE of less than 0.005. This search was repeated for each data point in the sample period and the resulting parameter values were averaged. Simulations were also run with parameters obtained from the averages of the values obtained from the search algorithm over sample periods including from 1 to11 data points. All sample periods were taken from the start of the data series.

In Figure 4, we show the actual and simulated tracks of the volume share of Brand A. In general, we would not expect the amplitude of the simulated track to be as great as the amplitude of the actual track because the simulation does not take account of special offers and the like.

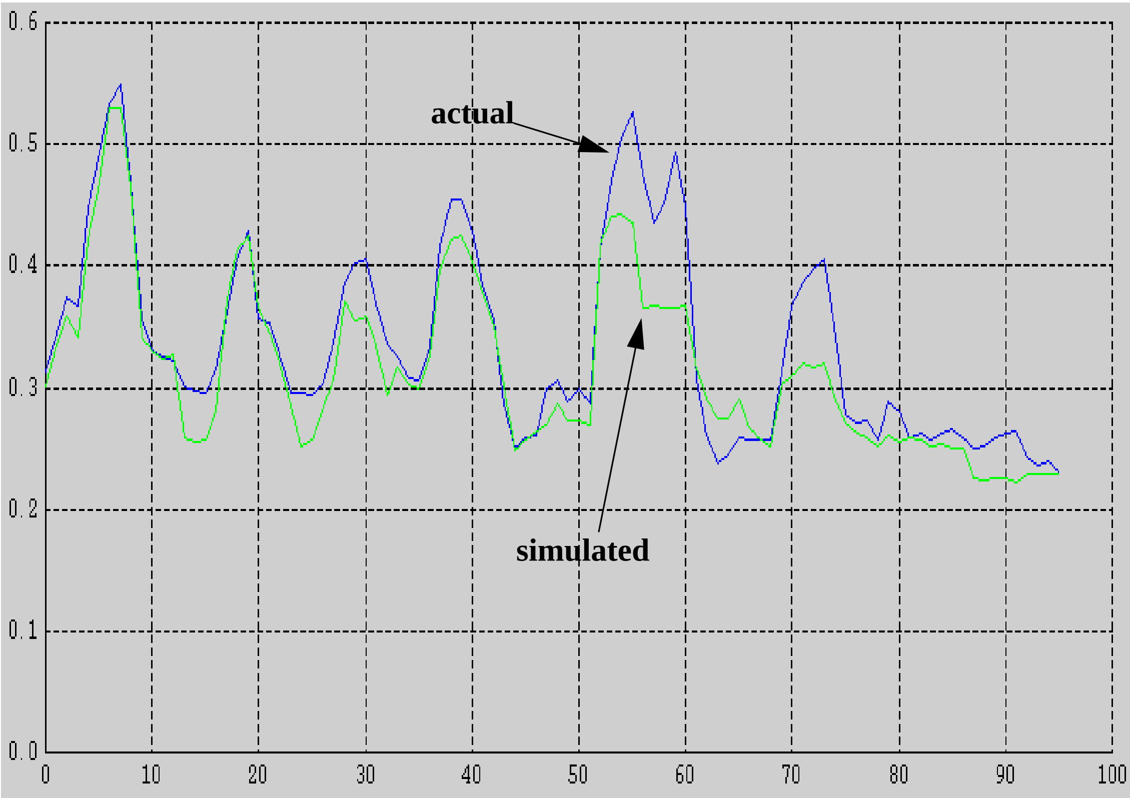

Figure 4: Actual and simulated market volume shares — Brand A (10-week sample period)

This simulated share track in Figure 5 is also close to the actual track and misses the largest fluctuations around Christmas. Unlike the track for Brand A, however, the Brand B simulated share track is everywhere above the actual track. Although we do not show them here, this bias is also reflected in the simulated share tracks for Brand Cand Brand D. The bias is certainly a result of a misspecification of the preference-states or their relative importance. The misspecification itself could be in the views expressed by the marketing professionals acting here as domain experts or in the parameterization of the model. In either event, the elimination of such bias is a matter for further investigation.

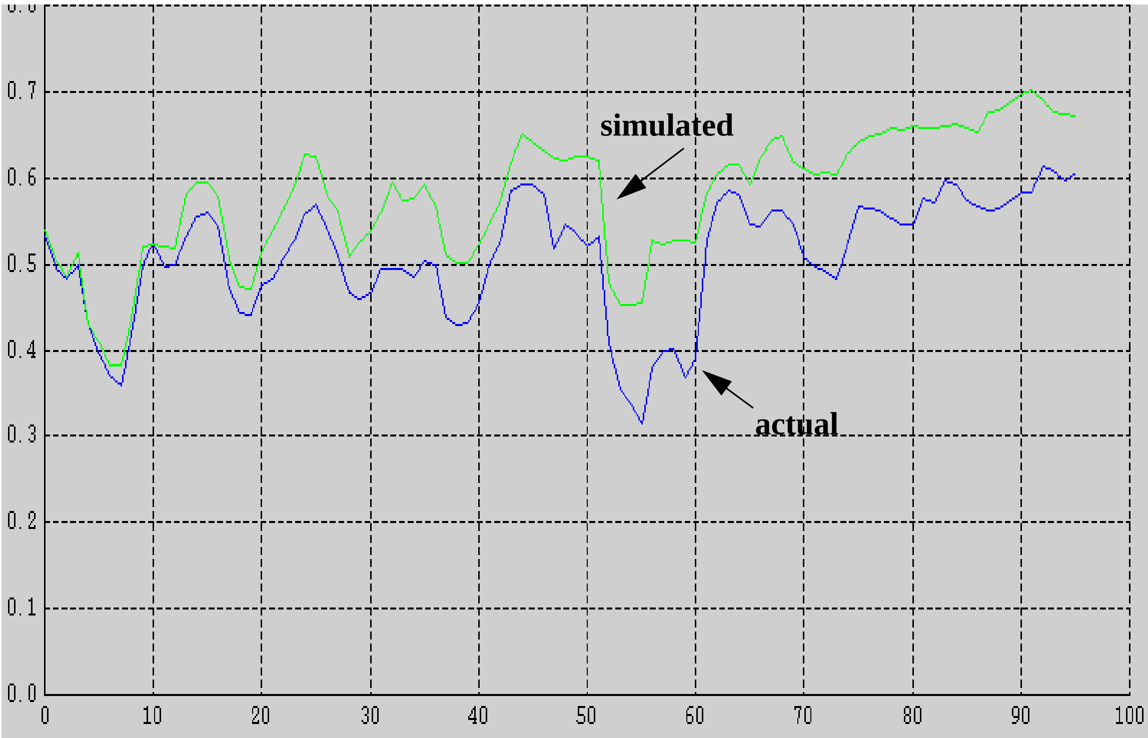

Figure 5: Actual and simulated market volume shares — Brand B

In order to get some more general and rigorous indication of the accuracy of the simulation, we regressed each of the volume shares against subsets of all four price series and then calculated the root mean squared errors and the mean absolute percentage errors for the shares generated by the estimated regression equations over the whole of the data set. We parameterized the preference-state model for 1, 3, 6 and 11 periods and averaged in order to simulate the remaining weeks of the data period. As shown in Table 3 and Table 4, the simulations of market shares with the preference-state model are unambiguously superior to the simulations using OLS-estimated models where only the first few data points are used in the parameterizations. The root mean squared errors (RMSEs)for all of the included shares were better for the preference-state model parameterized on the first data point only than for the OLS model estimated over the first 16 data points. Apart from the series for the Brand B, the single-data-point parameterization was also better in terms of RMSE than the OLS model estimated over the first 21 weeks' data.

Table 3: Comparing Preference-State and OLS models — RMSE (%)

	Brand A	Brand B	Brand C	Brand D
Preference-state model (1-week sample period)	5.4	9.8	2.4	3.9
Preference-state model (6-week sample period)	8.3	3.8	1.6	4.8
Preference-state model (11-week sample period)	3.6	7.8	2.1	3.8
OLS model (11-week sample period)	34.1	27.9	4.8	10.5
OLS model (16-week sample period)	24.8	9.8	6.4	9.3
OLS model (21-week sample period)	16.2	6.9	2.9	7.3
OLS model (51-week sample period)	3.5	2.8	1.2	1.5
OLS model (96-week sample period)	3.4	2.7	1.0	1.4

A similar result is observed with respect to mean absolute percentage errors (MAPEs) in Table 4. The results from parameterizing the model with the average of the parameter values obtained from minimizing absolute percentage errors for the first 11 data points yields MAPEs with the preference-state model which are better than those obtained with an OLS model estimated over the first 21 data points for every product except the own-label brands.

Our conclusion here must be that the preference-state model is much more expensive computationally than OLS models but is both more economical in its use of data and yields much more information. Although neither model is being used here for forecasting, it appears from the robustness of the simulation results that, for any given sample period, the preference-state model will yield more reliable scenarios for any given set of pricing data that would the OLS models. For the purposes of scenario development, these results give some confidence that the preference-state's implications for the market-share effects of different pricing strategies will not be worse and might well be substantially more accurate than the implications of the OLS-estimated models.

Table 4: Comparing Preference-State and OLS models — MAPE

	Brand A	Brand B	Brand C	Brand D
Preference-state model (1-week sample period)	11.8	18.3	41.4	31.8
Preference-state model (6-week sample period)	25.2	6.2	25.2	40.4
Preference-state model (11-week sample period)	7.2	13.9	36.1	30.7
OLS model (11-week sample period)	77.0	39.9	75.9	85.6
OLS model (16-week sample period)	73.3	17.9	113.4	76.9
OLS model (21-week sample period)	46.3	13.4	46.1	58.5
OLS model (51-week sample period)	14.5	8.2	19.2	15.2
OLS model (96-week sample period)	14.6	8.4	17.3	14.6

We turn now to the question of whether the data we used is in any way inappropriate for the comparison between OLS-estimated equations and the preference-state model.

In fact, the data seems entirely appropriate for OLS estimation as can be seen from the track of the Brand B share in Figure 6. This is typical of the four tracks obtained from the regression equations.

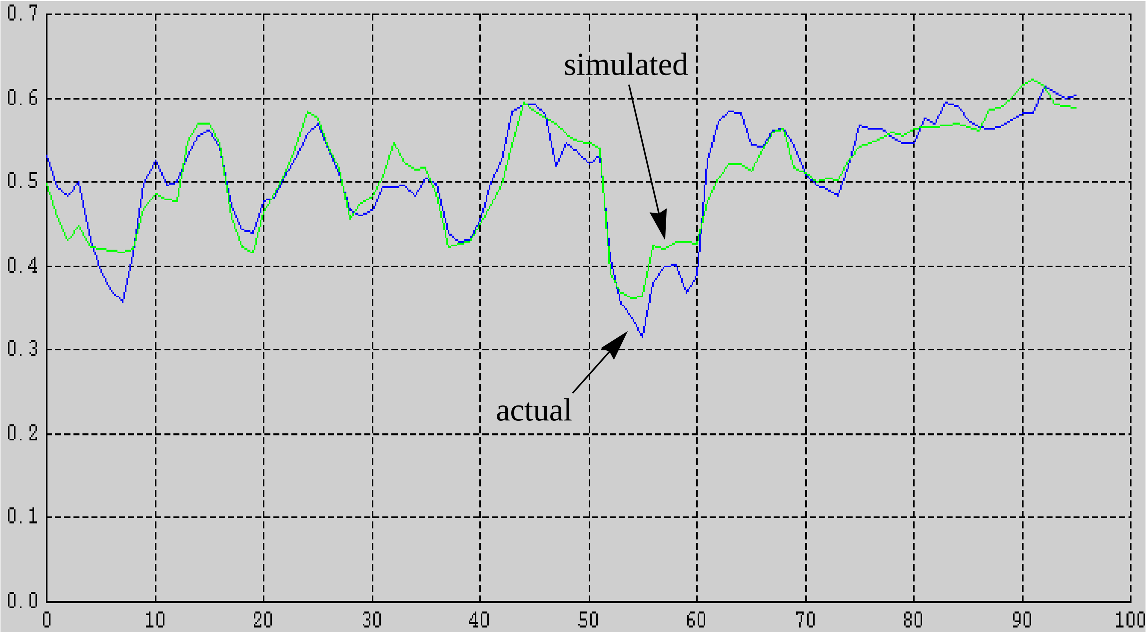

Figure 6: Brand B shares — actual and simulation using OLS regression equation estimated over whole data set

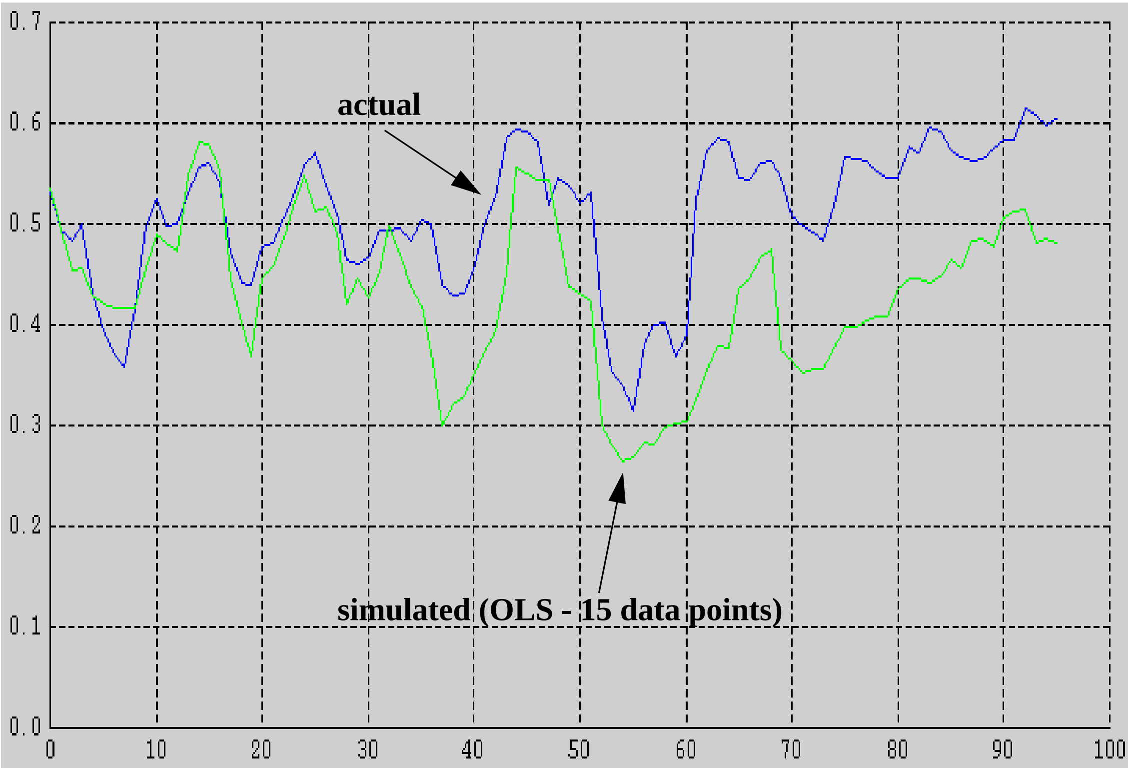

Figure 7: Brand B shares — actual and simulation using OLS regression equation estimated over weeks 0-14

However, the same regression equation for the Brand B share estimated on only the first 15 data points yields the simulated share track shown in Figure 7. That figure simply provides visual confirmation of the results reported above. The point here is only that the data set as a whole does not discriminate against the OLS-estimated model. While we cannot yet give a definitive reason for the goodness of the performance of the preference-state model, it fits with experience that estimation and forecasting in general are improved by domain knowledge. In this case, both the structure of the model and the specification of the qualitative variable values were based on extensive domain knowledge. The resulting model differed from standard models based on mainstream, economic preference theory in allowing for both non-convexities and asymmetries in the demand relations.

6 Conclusion

We have shown in this paper that models built on qualitative information elicited from domain experts yield numerical relationships which are not less robust than results obtained on suitable data by means of conventional statistical analyses. One feature of this approach is that the models contain a great deal more information that will be directly useful to domain experts (such as marketing professionals) than do statistical models.

We note that both OLS-estimated models and our preference-state model track market shares better in the earlier than in the later weeks of the observation period. An improved fit is obtained for the OLS-estimated model by fitting the earlier and later weeks' data separately. The fit of the preference-state model can be improved by changing, for example, the characteristics of the brands or the preferences attaching to the respective preference-states. Further regressions tell the analyst about changes in the sensitivity of individual shares to individual prices. Changing preference or characteristics parameters of the preference-state model imply judgements about qualitative changes in either demand patterns or consumer perceptions of brand characteristics. The model will not give a single, definitive answer but it will show whether those judgements are compatible with the numerical data and with the wider qualitative judgements of the marketing professionals.

The preference-state models thus provide the basis for a richer discourse in the analysis of sources of changing market-share patterns than do simple statistical models. And this added richness of discourse is not at the expense of formal modelling rectitude.

7 References

- Gordon, W. (1994), “Meeting the Challenge of Retail Brands”, *Admap* **29**(3), pp. 20-4.

- Morello, G. (1993), “The Hidden Dimensions of Marketing”, *Journal of the Market Research Society* **35**(4), pp. 293-313.

- 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.

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

- Horsky, D. and P. Nelson (1992), “New Brand Positioning and Pricing in an Oligopolistic Market” *Marketing Science* **11**(2), pp. 133-53.

- Deaton, A. and J. Muellbauer (1980), “An Almost Ideal Demand System”, *American Economic Review* **70**, pp. 312-26.

- Leitch and Wyatt, (1994)