Non-participation in choice models: hurdle and latent class models

Michael Burton and Dan Rigby[Ψ]

Paper presented at EnvEcon 2007
Friday 23rd March, Royal Society, London

Abstract

In repeated choice modelling studies, it is often the case that individuals always select the status quo option. Although this pattern may reflect considered choices, they may also be the result of alternative decisions about whether to participate in the choice process at all or. Alternative methods of dealing with this feature of such data, each with associated implications for estimates of economic values, are presented here. In particular we consider the alternative strategies of excluding such individuals from the data, using hurdle models to explicitly model this group, and propose the use of latent class models to endogenously allow for difference preference structures. An advantage of the latent class approach is that the form of the non-participation needs not be defined in advance. These approaches are considered using UK choice experiment data on food choices where the attributes include genetic modification of food. The data reveal two forms of non participation within the same model.

JEL: C8, D6, C23

Address for Correspondence:

Dr Dan Rigby
Economics
Manchester University
M13 9PL

Tel: +44 161 275 4808
Fax: +44 161 275 4812
email: dan.rigby@man.ac.uk

[Ψ] Rigby is Senior Lecturer, Economics, School of Social Sciences, University of Manchester. Burton is Associate Professor, Agricultural and Resource Economics, University of Western Australia. This paper draws work commissioned by UK Department of Environment Food and Rural Affairs (DEFRA). The views presented in this paper are those of the authors alone and should not be regarded as those of DEFRA or of individuals within DEFRA.

Non-participation in choice models: hurdle and latent class models

1. Introduction

While the use of stated preference studies in the agricultural and resource field proliferates, relatively little attention is paid to the levels and implications of respondents refusing to participate at all, or, if the survey design permits, only making choices that avoid certain characteristics of the good/scenarios. Within choice experiments such behaviour may be flagged by, for example, repeated selection of "Current" or "None of These" options. Typically, the statistical analysis of data generated by such surveys assumes that there are continuous utility functions which are underpinning choices. For those that have an extreme objection, that is not the case: they are making lexicographical choices. Analysis that ignores this possibility misrepresents the nature of preferences, and will provide inaccurate estimates of WTP and associated welfare effects.

In this paper the effects of two alternative approaches to incorporating these forms of preferences within choice experiment data are presented. First we follow von Haefen *et. al.* (2005) in using double hurdle models to explore whether the use of the mixed logit specification alters the results obtained when simply excluding those in the sample always choosing the current option. Essentially one is exploring whether some in the sample are really participating but have preferences which lie in the tail of the distribution (such as extreme aversion to an attribute). Unlike von Haefen *et. al.* we find the double hurdle models perform significantly better than a standard mixed logit model.

The novel approach presented here considers non-participation, and its effects on welfare estimates, using latent class models. This allows more varied and subtle forms of non-participation to be identified. It also permits their presence or otherwise to be statistically tested. While the hurdle model approach of von Haefen *et. al.* (2005) requires *a priori* definition of non-participation, this new approach does not. These two approaches are applied to choice experiment data from the UK in which attributes of bread are used to investigate choices and preferences. One of the attributes is the genetic modification of ingredients, an issue in the UK that is likely for some to involve lexicographic preferences. These may manifest themselves as repeated selection of the status quo option, but also consistent selection of non GM options.

The purpose of this paper is to apply two alternative approaches to incorporating lexicographic preferences within the analysis of choice experiment data. Their implications and potential for researchers are then discussed. The following section briefly outlines the standard methods of estimating choice models while Section 3 summarises the hurdle model approach to non-participation. Section 4 sets out the proposed latent class approach to the issue. The data and policy issue are then set out before results from both approaches are presented. Discussion of these results and conclusions follow.

2. Choice Experiments, Random Utility Theory and Preference Heterogeneity

Choice experiments have become common now in the environmental and resource economics literature. The technique is based on the idea that individuals can choose between alternative options comprising a number of attributes with different levels. Random utility theory proposes that individual consumers choose alternatives that yield the greatest utility and so the probability of selecting an alternative increases as the utility associated with it increases. Hence a person faces a choice among alternatives in choice set t on each of the occasions they make a choice. The utility that respondent n obtains from alternative j in choice situation t is:

$$U_{njt} = \beta' x_{njt} + \varepsilon_{njt} \qquad (1)$$

where x_{njt} is a vector of observed variables and coefficient vector ß, representing peoples' tastes.

The model is implemented by choosing a particular distribution of disturbances, typically the disturbances are assumed to be Gumbel distributed. This provides a variant of the logit model used in discrete choice modelling with the probability of person n choosing option j from N options (π_{nj}) expressed as:

$$\pi_{nj} = \frac{\exp(\beta' x_{nj})}{\sum_{j=1}^{N} \exp(\beta' x_{nj})} \qquad (2)$$

In this specification the scale factor has been normalised to one and the *t* subscript for the choice situation has been suppressed (see Louviere *et al* (2000)). In these models 'partworths' or Willingness to Pay values (WTPs) are obtained from the (negative) ratio of an attribute's marginal utility to the marginal utility of the payment vehicle.

There has been considerable effort applied to the better accommodation of differences in tastes within choice modelling survey data. In many studies demographic variables which may explain differences in choices and hence partworths are used to identify group-specific point estimates of WTP (e.g. Burton et al, 2001). An alternative and now common approach is to estimate a mixed or random parameter logit models (Train, 1998; Revelt and Train, 1998; Train, 1999) to identify the distribution from which preferences are drawn. These are a form of the random utility model in which it is assumed that the functional form and arguments of utility in (1) are common but that the ß 'taste' parameters vary across individuals.

In the mixed logit the coefficient vector in (1) is allowed to vary among the population with density f(ß|θ*). This vector of coefficients ($ß_n$) can be expressed as the population mean (b) and the individual specific deviation from that mean η_n. Hence the utility that respondent n obtains from alternative j in choice situation t is rewritten as:

$$U_{njt} = b' x_{njt} + \eta_n ' x_{njt} + \varepsilon_{njt} \quad (3)$$

Given that the values of $ß_n$ are not known, the probability of choosing option i in choice t is the integral of the conditional probability in (3) over all possible values of ß. This integral takes the form:

$$Q_{nit}(\theta^*) = \int L_{nit}(\beta) f(\beta \mid \theta^*) d\beta \quad (4)$$

Denoting the alternative that person n chose in period t as i(n,t) and assuming that $ß_n$ = ß, the probability of person n's observed sequence of choices is given by:

$$S_n(\beta) = \prod_t L_{ni(n,t)t}(\beta) \quad (5)$$

Given that $ß_n$ is unobserved, the probability for the sequence of choices is the integral of (5) evaluated over all possible values of ß which depends on the distribution of the ß:

$$P_n(\theta^*) = \int S_n(\beta) f(\beta \mid \theta^*) \; d\beta \quad (6)$$

The log-likelihood function is LL(θ)= Σ_n lnP_n (θ) which is maximised via simulation in which P_n (θ) is approximated by summing over values of ß generated by Halton draws (see Train, 1999). For a given value of the parameters θ, a value of ß is drawn from its distribution and on the basis of this draw S_n (ß), the product of standard logits, is calculated. The process is repeated and the mean of the values of S_n (ß) is interpreted as the estimate of the choice probability:

$$SP_n(\theta) = (1/R) \sum_{r=1,\dots,R} S_n(\beta^{r|\theta}) \quad (7)$$

where R is the number of draws of ß, $ß^{r|\theta}$ is the r^{th} draw, and SP_n (θ) is the simulated probability of person n's sequence of choices. The simulated log-likelihood function is SLL(θ) = Σ_n ln(SP_n (θ)) and the estimated parameters are those that maximize the function. A number of alternatives are feasible for the distribution of ß, including normal, log-normal, triangular and uniform.

3. Non-Participation in Choice Experiments

These random utility models have in common the assumption that there is a conventional utility function underlying the choices made, even if that utility function is individual specific, through the introduction of individual characteristics or random parameters into the model. Hurdle models change that assumption. Instead it is assumed that the sample comprises a mixture of two (or more) types, with different preferences for the good. The difference expresses itself in an extreme form: for some portion of the sample, the good under consideration does not enter into their relevant choice space: they simply do not consider it to be a valid choice of product and are hence non-participants in the market.

Within the context of choice modelling, one has to be clear as to the definition of non-participation. von Haefen *et. al.* (2005) use consistent selection of the *status quo* option within a series of choice estimates as an indication that the respondent is a non-participant. Adamowicz et al (1998) suggest that repeated selection of options that have the lowest (or highest) value for a particular attribute may also be an indication of non-participation (i.e. a simple heuristic is being used which avoids the need to make considered judgments across all attributes and levels).

Assuming there are T choice occasions, the probability that an individual using the conventional utility function in (1) will repeatedly select the status quo (option 1) in all of them is given by:

$$\prod_{t}^{T} \pi_{n1}^{t} \tag{8}$$

where π is the probability given by (2). The issue is that the observed number of individuals who behave in this way is often high, given the number of choice occasions they are facing (For example, James and Burton e.g. 31% of the sample, faced with 9 choice sets always selected the status quo option in James and Burton (2003). The alternative is to assume that there is an independent probability that an individual is a non-participant, $\bar{\pi}$. The likelihood that an individual will generate a 'non-participation' outcome is now the sum of the two:

$$\bar{\pi} + \prod_{t}^{T} \pi_{n1}^{t} \qquad (9)$$

Thus, although the double hurdle model introduces a separate probability of non-participation, it assumes that it is still possible that some element of non-participation is generated by the conventional discrete choice framework. This introduces the possibility that the observed serial selection of the *status quo* is due to one of two causes: genuine non-participation in the process or genuine selection of serial *status quo* as a corner solution. Empirically it is not possible to differentiate between these two causes from simple observation of the data, but it can be accommodated within the estimation of the model by reforming the likelihood function for the model (see von Haefen *et. al.* 2005). This approach is more intuitively appealing when using a mixed logit model for the conventional utility function where the distribution of preferences may be such that preferences in the tail of the distribution may lead to choice patterns which are identified as non-participation but are in fact the result of weighing up attributes rationally and hence participating.

It is also possible to make the non-participation probability $\bar{\pi}$ a (suitably defined) function of attributes, so that the causes of non-participation can be examined. In the application below we specify it with a probit function. There is no reason why the use of the hurdle approach is limited to the fixed parameter discrete choice model: it is possible to introduce combined mixture models where there are random parameters associated with both the non-participation and utility functions.

4. A Latent Class Approach to Non-Participation

An alternative approach which resembles mixed logit but which relaxes the requirement to make specific assumptions about the distributions of parameters across individual consumers is the latent class (LC) model (Boxall and Adamowicz, 2002). In this model consumers are assumed to belong to different segments or classes, each of which is characterised by unique class-specific utility parameters. A consumer can probabilistically be assigned to one of several latent classes depending on his or her characteristics and preferences. In the LC model the probability that person n will choose alternative j is defined as follows:

$$\pi_{nj}^{'} = \sum_{s=1}^{S} \pi_{nj \bullet s} \pi_{ns} \qquad (10)$$

where s denotes the number of segments or classes, $\pi_{nj \bullet s}$ is the probability that individual n chooses alternative j conditional on class s, and π_{ns} is the marginal probability that individual n is in class s. This can be viewed as a generalisation of the double hurdle model, with a number of underlying utility functions and associated choice probabilities $\pi_{nj \bullet s}$ for different classes, while non-participation is not defined *ex ante* by an observed form of choice behaviour, but is determined by special forms of the utility function for some classes, and the (endogenous) probability of being a member of that class. The latter can then be expressed as:

$$\pi_{ns} = \frac{\exp(\gamma'_s C_n)}{\sum_{s=1}^{S} \exp(\gamma'_s C_n)} \qquad (11)$$

where γ_s denote a set of class-specific coefficients on concomitant variables C_n.

If one assumes that one of the classes is one of non-participation one can construct prior expectations of parameter values within the utility function for that class. Thus, if a *status quo* alternative specific constant is employed in the specification, then one would expect a high and limiting value for that parameter, ensuring that all those within the class have a very high probability of selecting that option throughout the choice sequence, regardless of other attribute values. Indeed that class may be characterised by insignificant parameters on the other attributes, including price. The advantage of the approach is that it does not require *a priori* imposition of a definition of non-participation, and potentially will allow one to identify a number of limiting behaviours within the data. For example, it may allow one to simultaneously identify sub-populations who, in this study, select only the status quo option, always avoid a certain attribute (level), or only use price as a basis for choice.

Having outlined the issue of non-participation in choice studies we now describe the application in terms of the substantive policy issue and the CM survey from which the data which is used to explore both latent class and hurdle models of non-participation.

5. Lexicographic Preferences and Non-Participation: GM Food in the UK

A number of studies of attitudes in the UK and mainland Europe have found that many consumers do not want to eat GM food and that the majority believe that, if such food is sold, it should be clearly labelled (see Consumer Association, 2003; MORI, 2003; Marris *et al*., 2001; various Eurobarometer surveys; the *GM Nation?* consultation). As part of its attempt to defuse the WTO dispute with US, Canada and Argentina whilst responding to domestic pressures, the EU introduced a new labelling regime in 2004 (EU Regulations 1829/2003, 1830/2003). This extended the labelling regime to include not only those foods containing modified genetic material (which we term 'GM Food') but also those foods with ingredients derived from GMOs, such as maize and soy oil, despite the absence of modified DNA or protein in them (we term this 'GM Derived Food'). GM Derived foods, which may be indistinguishable in laboratory tests from non-GM counterparts, now required a label. The detectability of genetically modified DNA or protein in the final food product was no longer the basis of labelling, an extension to process based labelling.

There are compliance costs associated with this change which can only be justified if consumers place a value on the information conveyed by the GM Derived label. The benefits (if any) of this change to the labelling regime were explored via a choice modelling survey in the UK in which respondents were asked to make choices among varieties of bread comprising attributes including the GM nature (or otherwise) of the ingredients.

The survey was conducted in England, Wales and Scotland between July and September 2003. The sample was defined as men and women, aged 16 and over who were the main food shopper for their household. The choice set attributes and levels were finalised following a series of semi-structured interviews and a number of pilot interviews. The bread attributes and levels concerning the price, fibre content and life of the bread before becoming stale as well as the GM, GM Derived (GMD) or Non-GM nature of the ingredients are reported in Table 1. The sample employed a main effects design and a sample comprising 608 respondents was achieved. Personal interviews were conducted in the home using computer aided personal interviews with each respondent was required to answer 4 choice modelling sets, following some (discarded) practice sets. The first option in each set was identified as the status quo i.e. their current bread which in all cases was non-GM.

274 (45%) of the 608 respondents selected the status quo in all 4 choice sets. However, the inclusion of some non-GM breads alongside the SQ option introduces an alternative measure of non-participation: the possibility of dismissing any options that contain GM, but making choices across the non-GM breads offered. 157 respondents (26%) were in this category: they made some

non-status quo selections, but never selected a bread. With GM The remaining portion of the sample made some selection involving GM or GMD bread across the 4 choice sets presented. It is to the implications and decomposition of these behaviours that we now turn.

6. Results: Hurdle Models

We first address the participation issue using hurdle models before the proposed use of latent class models to consider this issue is outlined. It is assumed that the utility function of those who participate contains the 4 attributes of the choice sets, with Price, Shelflife and Fibre content included as continuous variables. The nature of the GM technology used is included as 2 dummy variables, for GM Derived and GM, using non-GM as the baseline. Gender and age effects are interacted with the GM term. An alternative specific constant (ASC) is included for the status quo option. Von Haefen et al (2005) found that the use of a random parameter representation of the utility function significantly altered the inferences about the value of the double hurdle model for their data set. Specifically, the double hurdle model was preferred over the conventional conditional logit model when parameters were assumed fixed, but showed no increase in explanatory power when a random parameter model was used. The rationale for this is that the greater flexibility in representing preference heterogeneity inherent in the random parameter model means that there is no requirement for the additional 'mixture' of non-participation. Because of this result, we explore the consequences of introducing a random parameter on the ASC variable given the potential importance of this variable in explaining non-participation.

The use of the hurdle model requires definition of what is non-participation and here we define it as selection of the status quo option in all 4 choice sets. The probability that an individual is in the non-participation group is modelled as a probit, with a number of individual specific variables, including age, membership of social class and a PCA composite attitudinal score regarding 'trust' in GM food. It is important to note that in the double hurdle model the probit and mixed logit model are estimated jointly, as one has to account for the probability that an observed 'non-participant' is in fact a genuine participant who selects the status quo in each choice set.

Four models are estimated[1] for the full data set: conditional logit with and without a random parameter for the 'Current' ASC (denoted Conditional Logit [CL] and Mixed Logit [ML] respectively) and a double hurdle model with and without a random parameter estimate for the ASC (denoted Double Hurdle Conditional Logit [DHCL] and Double Hurdle Mixed Logit [DHML] respectively). Because the conventional and double hurdle models are not nested, a non-parametric variant of the Vuong test is used to compare them (Clark, 2003). This is a distribution-free test that applies a modified paired sign test to the differences in the individual log-likelihoods from two non-nested models. Table 2 reports estimates for pairwise comparisons of all 4 models. It is clear that the Conditional Logit is dominated by the three other models (ML, DHCL, DHML), in which heterogeneity is represented either by a random parameter specification, a double hurdle model, or both. The test results do not allow one to discriminate between the ML, and DHCL models. However when comparing random parameter models, the DHML model dominates the ML model and the DHCL model, leading us to prefer the double hurdle mixed logit specification. This is in contrast to the results in von Haefen et al (2005) where introducing either a hurdle or random parameters represented a significant improvement in model fit, but it was not possible to discriminate between the two.

As our interest is in the implications of the double hurdle model, in the discussion that follows we compare only Mixed Logit and Double Hurdle Mixed Logit models. Table 3 reports the parameter

[1] Estimation of the double hurdle model employed the GAUSS code provided by Roger von Haefen: assistance in implementing the code is gratefully acknowledged.

estimates of the Double Hurdle Mixed Logit and Mixed Logit models, plus an additional model for comparison.

Model 1 reports the Mixed Logit estimated over the full data set, including those identified as non-participants (i.e. the ML model reported in Table 2). Model 2 reports the estimates from a Double Hurdle mixed logit model applied to the full data set. Three variables are used to explain non-participation: age, being a member of social class E (state pensioners, casual or lowest grade workers), and the degree of 'trust' in the use of GM food. This PCA trust term was derived from Likert scale responses to a series of statements regarding food biotechnology (see Rigby *et al*., 2004). This attitudinal term is significant, and implies that those who are more positive regarding GM technology, its use in the food system and its regulation are less likely to be among the non-participant group.

Model 3 is provided as a further comparison: it is estimated using a Mixed Logit model on the subset of the data that excludes those identified as non-participants. This is the method most often used to accommodate non-participation, but it assumes seperability i.e. that there is no loss in efficiency by imposing an *a priori* segregation of the data. As von Haefen *et. al.* (2005) note, this will deliver the true parameters only if one can accommodate for truncation of the data i.e. the impossibility of repeated selection of the *status quo*. The model identifies significant price and GM and GMD effects, with increasing concern about GM as people get older (although the quadratic term implies this increasing disutility reaches a turning point at approximately age 50). There is no significant impact of gender on preferences for GM Derived food, and only at the 10% level for GM food with women's utility reduced more than men's by the presence of GM ingredients in their bread.. The 'status quo' ASC is estimated to have a negative mean but there is a substantial variance around this mean denoted by the ASC standard deviation (SQ sd) term.

Although direct comparison of parameters across these estimated models is strictly inappropriate, because of the confounding impact of the scale parameter, casual inspection of the results indicates that they are very similar in their coefficient estimates, apart from the status quo (SQ) ASC. As one may expect, the ASC is positive (with a significant mean) in the ML model estimated over the full data set, reflecting the substantial portion of the sample who always select this option, and there is a highly significant, and large, standard deviation for the distribution of the ASC, suggesting a greater variability in the sample compared with the restricted sample, which excludes "non-participants", which one would expect.

Within the mixed logit component of the DHML model, the parameter estimates are very similar to those in Models 1, suggesting that both models are giving equivalent representations of the underlying utility function, with different statistical representations of the sub-sample who are non-participants.

Figure 1 reports a scatter graph of the parameters from the three models against each other. There is a strong linear relationship between them, with the only notable difference being those for the mean and variance of the distribution of status quo ASC preferences, confirming that the treatment of non-participation in this case is not affecting the estimates of the relative weights of the attributes within the utility function.

7. Results: Latent Class Models

Respondents' choices and the definition and extent of serial non-participation were then reconsidered using latent class models, estimated using the full dataset. Table 4 shows the results from such a model (Model 4) in which there are 3 latent classes. A number of models with alternative numbers of classes were estimated but indicative statistics such as the log likelihood and AIC and BIC scores all suggested the 3 class model was preferred. Note that within this model no *a priori* assumption is made about participation or non-participation. Demographic variables are excluded from the conditional logit model, but are included in the model explaining class membership probabilities.

For Class 1 the parameter on the status quo variable is estimated to be positive and those for the GM and GM Derived terms variables are very large and negative which, for almost any values of the other attributes, leads to an extremely high probability of selecting the status quo option. Class 2 has very strong negative coefficients for the GM variables, implying that within this class, the presence of GM elements is strongly objected to, but the other attributes are significant apart from the status quo variable, which is insignificant. This suggests that this is a class of individuals who are averse to GM, but are prepared to make choices across the other attributes of the choice sets if they have the opportunity to compare breads with no GM or GMD ingredients. Class 3 has an insignificant GM Derived coefficient, but a significant and negative effect of GM ingredients. For this group the presence or otherwise of GM Derived ingredients in the bread is irrelevant. They are prepared to evaluate the breads on the basis of the more conventional attributes, however, they are averse to the presence of GM ingredients, but not to the extent that its presence is a limiting factor in choices (a point discussed further below).

We next test whether these results can be restricted so as to construe Class 1 as representing a form of non-participation. This is done via a series of Likelihood ratio tests. This 'Restricted Model'' also shown in Table 4, has restrictions on Class 1 utility function coefficients tested. Specifically, we test whether it is possible to restrict the utility function of Class 1 to one comprising just SQ and GM effects, removing the role of Price, Fibre and Shelflife in determining choices. This restriction is accepted[2] (test statistic of 2.4, $\chi^2_{0.05,3} = 7.82$). The parameter estimates for the utility functions for Classes 2 and 3 and the class membership terms are stable across both model specifications, indicating the restrictions have not been accepted via the radical restructuring the nature of the classes nor the determinants of class membership.

The determinants of the class membership have a number of significant coefficients. The PCA composite attitudinal score *GMtrust* was a significant determinant of class membership, as were two additional PCA attitudinal scores '*Green'* and '*Active'*[3]. Those who are more trustful of GM food and its regulation are less likely to be members of Classes 1 and 2, compared to Class 3. A positive score on the *Green* PCA score increased the chance of being in Class 2 and also, but to a lesser extent, Class1 relative to Class 3. The effect of higher *Active* scores was positive and equivalent for both classes 1 and 2. Age also plays a role in explaining class membership probability. As age increased so did the probability of being in the non-participation Class 1, while there was a marginally significant negative effect on the probability of being in Class 2.

Considering the sign, magnitude and significance of the utility function coefficients of the 3 classes, the 3 classes could be characterised in the following way. Class 1: 'Non Participators', Class 2: 'Participating but GM Averse', Class 3: 'Participating'. This characterisation is examined by comparing the individual- specific predicted class membership probabilities and the observed choice behaviour of those individuals. Specifically, the sample is split into 3 groups on the basis of their choices. The first group always chose the SQ option (45% sample), the second group selected one or more non-SQ options but never a GM bread (26%) and the third group selected one or more

---

[2] We thank William Greene for modifying the NLOGIT code to allow these restrictions to be imposed.

[3] The '*Green'* score is based on Likert responses to the statements: "I try to avoid artificial ingredients", "I try to recycle as much waste as possible", "When I have the choice I always buy organic", "I try to buy environmentally friendly products" and " When I have the choice, I always try and buy ethically responsible products (e.g. Fair Trade)"). The PCA composite 'Active' was derived from responses to the following statements: 'Food should be clearly labelled to say if it contains genetically modified (GM) ingredients', 'I read ingredients labels on food items very carefully', 'I consider myself to be knowledgeable about food safety issues', 'I don't mind paying extra for quality food'.

GM bread choices (29%). The mean predicted class membership probabilities for these 3 observed groups are shown in Table 5.

The 3 classes shown in the restricted model in Table 4 fit the 3 characterisations of observed behaviour in the data well. Class 1 representing a form of non-participation, Classes 2 and 3 representing forms of participation. Class 2 members appear to exhibit a hierarchical structure to choices, with the presence of GM in any form leading to rejection of an option, but the other attributes being used as a basis for choices between breads that do not include GM. The overall mean membership probabilities also closely approximate the frequency in the overall sample of the 3 types of choice behaviour. We expect this matching of behaviour types to classes to cause the 3 class model to be preferred over models with alternative numbers of classes.

Finally, we consider the *conditional partworths* from the latent class models, estimated as conditional upon being in a particular class (Table 6). Although it is more usual to report distributions of unconditional partworths from latent class models, in the current case this is infeasible, as in Model 5 the utility function for Class 1 has a zero coefficient on price, which generates irrelevant estimates of partworths. Note that these partworths are in % terms, where 100% represents the cost of a loaf of bread which in this sample averaged close to 1€/$1.25. The results reveal that, as the discussion above indicated, those in Class 2 are strongly averse to the GM aspects of the bread, but are relatively unresponsive to changes in the levels of the other attributes. This contrasts with member of Class 3, which appears to be indifferent towards GM Derived bread, concerned about GM but less so than members of the other classes, hence the WTP of €0.21 to avoid GM bread. The valuations of changes in the Fibre and Shelflife attributes for Class 3 are also higher than the other 2 segments of the market.

8. Discussion and Conclusions

This paper has explored two alternative methods of accounting for non-participation in choice experiments. Such behaviour, typically identified as serial selection of the 'status quo' or 'none of these' options typically results in (significant) numbers of respondents' choices being excluded from the data and subsequent analysis. The frequency of this non-participating behaviour is likely to increase where one or more of the attributes is regarded as controversial and therefore likely to invoke lexicographic responses. GM food, in much of Europe, is such an issue and hence non-participation is analysed using choice modelling data concerning this foodtype.

We follow von Haefen *et. al.* (2005) in using double hurdle models to explore whether the use of the mixed logit specification alters the results obtained when simply excluding those in the sample always choosing the SQ option. Essentially one is exploring whether some in the sample are really participating but have preferences which lie in the tail of the distribution (e.g. extreme GM aversion).

The double hurdle model, where non-participation is defined as serial selection of the status quo option, performed significantly better than a standard mixed logit model. However it changed little in the way of relative weights of the attributes, apart from the alternative specific constant for the status quo option. Furthermore, this conclusion also held when comparing results with a mixed logit model applied only to those identified as 'participants' implying that one can impose seperability between the participation and choice components of this data set. This suggests that there is relatively little lost in terms of understanding behaviour of those making choices from taking the conventional approach to non-participation, which is to exclude them from the data set, but it should be emphasised that this is likely to be a study specific result, and may not hold for other data sets.

An alternative approach for analysing serial non-participation using a latent class model is proposed. This is achieved by imposing and testing restrictions that utility function coefficients for certain attributes in certain classes are insignificantly different from zero. This approach does not require the *a priori* definition of non-participation, but allows the data to reveal such behaviour(s).

Alternative specifications of the latent class model indicate the 3 class model is preferred. Analysis of the results from the model in conjunction with observed behaviour suggests 3 distinct forms of choice behaviour. The first type is non-participation as typically conceived, that is, serial SQ selection. The second type of behaviour suggests participation is occurring in that one or more non-SQ options is being selected, but never a GM option. This suggests the trading off over some attributes is occurring but not against the GM option (unless the range of attribute values is not sufficient to induce switching to a GM choice). Finally the third type of behaviour is characterised as the trading off across all attributes.

While past studies have hypothesised about the nature of latent classes, the first use here of formal restrictions on parameters within the conditional logit element of the latent class model is particularly useful in allowing the testing of prior hypotheses about the redundancy of some attributes for certain classes.

In this case three forms of response to GM content of bread is identified: those who appear to respond to the experimental structure of the study by only selecting the status quo option as a heuristic device to avoid GM; those who actively avoid GM through making choices across the options presented to them; and those who are indifferent to GM derived ingredients, and have some aversion to GM content, but are prepared to trade that aversion with other attributes. The motivations of the first group would be an interesting area for further study: whether this is a form of protest against the repeated choice experiment structure itself, or a simplifying heuristic to avoid GM. However, the remaining two classes imply that there is heterogeneity within this population with respect to GM, and in particular to GM derived ingredients: some find them completely equivalent to conventional ingredients while others find them completely unacceptable. This suggests that labelling GM derived may have a role in meeting some considered consumer preferences.

References

Adamowicz, W.L., P.C.Boxall, M.Williams and J.J.Louviere (1998) Stated preference approaches for measuring passive use values: choice experiments and contingent valuation. *American Journal of Agricultural Economics* 80: 64-75

Boxall, P.C. and W.L. Adamowicz (2002) Understanding heterogeneous preferences In random utility models: a latent class approach. *Environment and Resource Economics* 23: 421-446.

Burton, M.P., D.Rigby, T.Young and S.James (2001) Consumer attitudes to genetically modified organisms in food in the UK *European Review of Agricultural Economics*, Vol28(4):479-498

Clark, K.A. (2003) Nonparametric model discrimination in international relations. *Journal of Conflict Resolution* vol 47(1), 72-93.

Consumers Association (2002) GM Dilemmas - Consumers and Genetically Modified Foods. Policy Report http://www.which.net/campaigns/food/gm/misc/gm_report.pdf

von Haefen, Roger H, D. M. Massey, & W.Adamowicz. (2005) Serial Non-Participation in Repeated Discrete Choice Models, *American Journal of Agricultural Economics*. vol. 87(4) 1061-1076

Hensher, D.A. and W.H. Greene (2003) The Mixed Logit: the State of Practice, *Transportation* 30(2) 133-176.

James, S and Burton, M.P. (2003) Consumer preferences for food attributes with a focus on gene technology. *Australian Journal of Agricultural and Resource Economics* Vol.47:4 99501-518

Louviere, J. J., D. A. Hensher and J. Swait (2000). *Stated Choice Methods: Analysis and Application*, Cambridge University Press

Marris, C, Wynne, B, Simmons, P and S.Weldon (2001) *Public Perceptions of Agricultural Biotechnologies in Europe.* Final Report of the PABE research project. http://www.pabe.net

MORI (2003) "GM Food Opposition Continues" 3 July 2003 http://www.mori.com/polls/2003/gmfood.shtml

Revelt, D. and K. Train (1998), 'Mixed logit with repeated choices', *Review of Economics and Statistics* (80): 647–657.

Rigby, D, Young, T and M.Burton (2004) *Consumer Willingness to pay to reduce GMOs in food and increase the robustness of GM labelling*. Report to UK Dept of Environment, Food and Rural Affairs (DEFRA).

Train, K (1998) Recreation Demand Models with Taste Differences Over People. *Land Economics*, Vol. 74(2): 230-239.

Train, K (1999) *Halton Sequences for Mixed Logit.* Working paper, University of California, Department of Economics.

Figure 1. Scatterplot of attribute parameter estimates from Models 1, 2 and 3.

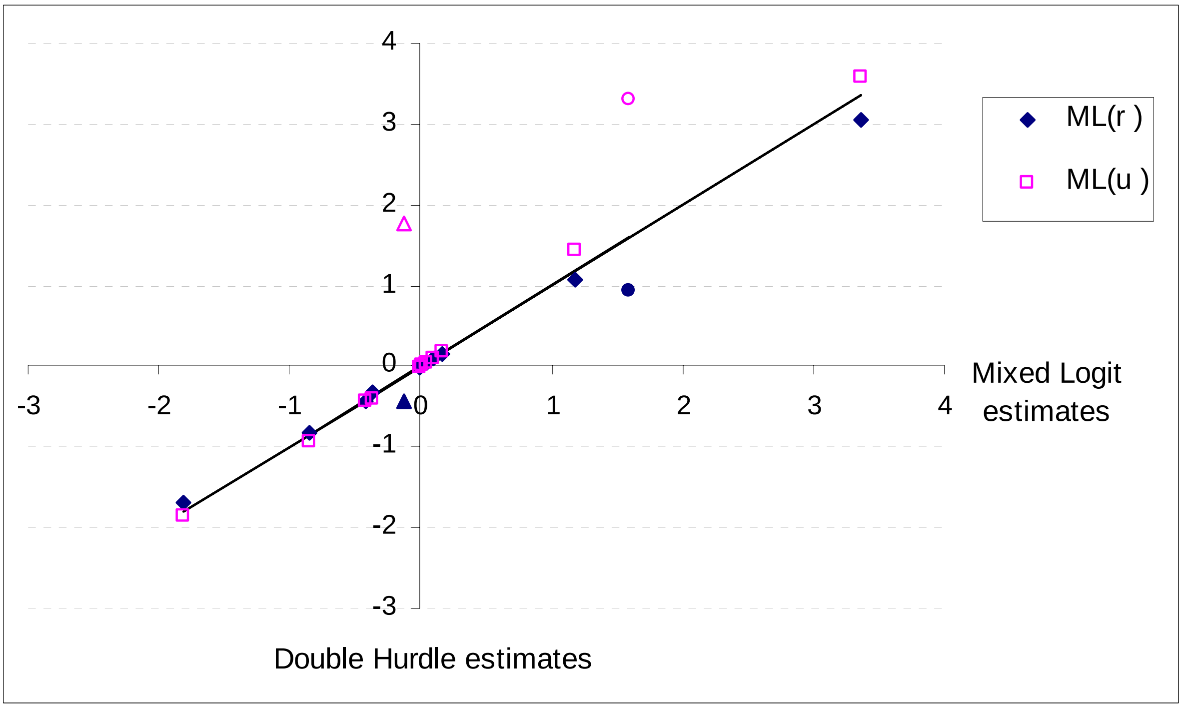

Note: SQ coefficients, mean and st.dev. denoted by triangles and circle respectively.

Table 1. Attributes and levels in the CM design.

Attribute	Levels
Price (%)	-67, -50, -33, -17, Usual, +17, +33
GM Type	Non-GM, GM-Derived, GM
Shelflife	Usual, Usual + 1 day, Usual + 2 days, Usual + 3 days
Fibre Content	Usual, Usual + 10%, Usual + 30%, Usual + 50%

Table 2 Tests for alternative specifications

		DH ML		DH CL		ML	CL
	LL values	*-0.687429*		*-0.700134*		*-0.696013*	*-0.847312*
DH ML							
DH CL	↑	χ^2 (<0.000)					
ML	↑	V (0.019)	=	V(0.979)			
CL	↑	V (<0.000)	↑	(<0.000)	↑	χ^2(<0.000)	

Notes:
Preferred model indicated by ↑; = indicates models cannot be differentiated between.
χ^2 :model comparison using standard χ^2 test; V: model comparison using Clark variant of the Vuong test
Significance level reported in parentheses

Table 3 Results from Mixed Logit and Hurdle Models

	Model 1 Mixed Logit: Unrestricted data set			Model 2 Double hurdle Mixed Logit			Model 3 Mixed Logit: restricted data set		
	parameter		*s.e.*	parameter		*s.e.*	parameter		*s.e.*
Hurdle									
const				-0.8691	***	*0.1958*			
gmtrust				-0.1802	***	*0.0623*			
age				0.0107	***	*0.0035*			
soc class E				0.3201	*	*0.1898*			
Discrete choice									
price	-0.0138	***	*0.002*	-0.0134	***	*0.002*	-0.013	***	*0.0019*
SQ	1.7646	***	*0.201*	-0.1281		*0.1768*	-0.4339	***	*0.1142*
GMD	3.5936	***	*1.0507*	3.3574	***	*1.0263*	3.0613	**	*0.9869*
GM	1.4229		*0.996*	1.1755		*0.9895*	1.0782		*0.9526*
GMD-fem	-0.4009		*0.3032*	-0.3667		*0.2939*	-0.3349		*0.2785*
GM-fem	-0.4397		*0.2698*	-0.4284		*0.2702*	-0.4374	*	*0.261*
GMD-age	-1.8609	***	*0.4281*	-1.8068	***	*0.4192*	-1.6925	***	*0.4023*
GM-age	-0.9274	**	*0.3762*	-0.8587	**	*0.3751*	-0.8135	**	*0.3628*
GM-age^2	0.174	***	*0.0424*	0.1724	***	*0.0416*	0.1619	***	*0.0398*
GMD-age^2	0.0927	***	*0.0358*	0.0889	**	*0.0359*	0.0848	**	*0.0349*
shelf	0.0526		*0.0421*	0.0508		*0.0419*	0.0481		*0.0414*
fibre	0.0037		*0.0029*	0.0038		*0.0029*	0.0041		*0.0028*
SQ st.dev	3.3165	***	*0.2077*	1.5967	***	0.2246	0.9246	***	*0.0969*

Note: Restricted dataset comprises those defined as participants. Unrestricted dataset comprises all individuals, both participants and non-participants.
***, **, * indicate significance at 1%, 5%, and 10% respectively.

Table 4 Parameter estimates for the latent class model

	Model 4: Unrestricted model			Model 5: Restricted model		
	parameter		*s.e.*	parameter		*s.e.*
	Utility function parameters: class 1					
SQ	2.213	***	*0.544*	3.083	***	*0.220*
Price	0.013		*0.010*			
GM	-2.894	***	*0.838*	-2.695	***	*0.828*
GM Derived	-2.252	***	*0.684*	-2.084	***	*0.619*
Shelf	-0.113		*0.222*			
Fibre	-0.020	**	*0.012*			
	Utility function parameters: class 2					
SQ	0.061		*0.284*	0.021		*0.289*
Price	-0.038	***	*0.006*	-0.039	***	*0.006*
GM	-5.773	***	*0.527*	-5.863	***	*0.542*
GM Derived	-5.334	***	*0.441*	-5.434	***	*0.455*
Shelf	0.203	***	*0.098*	0.216	***	*0.100*
Fibre	0.013	***	*0.006*	0.012	***	*0.006*
	Utility function parameters: class 3					
SQ	-0.614	***	*0.147*	-0.619	***	*0.146*
Price	-0.013	***	*0.002*	-0.012	***	*0.002*
GM	-0.250	***	*0.113*	-0.249	***	*0.113*
GM Derived	0.061		*0.113*	0.049		*0.113*
Shelf	0.154	***	*0.038*	0.152	***	*0.038*
Fibre	0.010	***	*0.003*	0.010	***	*0.003*
	Class membership parameters: class 1					
Constant	0.387		*0.338*	0.396		*0.337*
PCA: Gmtrust	-0.824	***	*0.137*	-0.823	***	*0.137*
PCA: Green	0.187		*0.128*	0.184		*0.128*
PCA: Active	0.264	***	*0.122*	0.266	***	*0.122*
Age	0.012	**	*0.006*	0.012	**	*0.006*
	Class membership parameters: class 2					
Constant	0.599		*0.439*	0.590		*0.442*
PCA: Gmtrust	-1.132	***	*0.176*	-1.138	***	*0.177*
PCA: Green	0.373	***	*0.183*	0.377	***	*0.185*
PCA: Active	0.253		*0.171*	0.254		*0.172*
Age	-0.013		*0.009*	-0.014		*0.009*
LL value	-1535.1			-1536.8		

Table 5. Observed Behaviour and Mean Latent Class Probabilities

	Mean Class Membership Probabilities (%)		
	Class 1	Class 2	Class 3
Observed Behaviour:			
Always SQ	99.4	0.6	0.0
Never GM, >1 non-SQ choice	24.9	73.2	1.8
>1 GM choice	3.6	7.4	88.9
All	52.3	21.3	26.3

Table 6 Conditional WTP for unit changes in attributes from the latent class models: % change in price of bread (100% = 1€).

	Unrestricted model			Restricted model		
	Class 1	Class 2	Class 3	Class 1	Class 2	Class 3
SQ	-170.2	1.6	-47.2	N/a	0.5	-51.6
GM	222.6	-151.9	-19.2	N/a	-150.3	-20.8
GM Derived	173.2	-140.4	4.7	N/a	-139.3	4.1
Shelf	8.7	5.3	11.8	N/a	5.5	12.7
Fibre	1.5	0.3	0.8	N/a	0.3	0.8