



THE MICROSIMULATION
UNIT

Number MU0101
DAE Working Papers
Amalgamated Series No. 0110

Combining household income and expenditure data in policy simulations

Holly Sutherland, Rebecca Taylor
and Joanna Gomulka

January 2001



Department of Applied Economics
University of Cambridge
Sidgwick Avenue
Cambridge CB3 9DE

Combining household income and expenditure data in policy simulations

Holly Sutherland, Rebecca Taylor and Joanna Gomulka

Microsimulation Unit Discussion Paper MU0101

January 2001

ABSTRACT

The analysis of the distributional impact of fiscal policy proposals often requires information on household expenditures and incomes. It is unusual to have one data source with high quality information on both and this problem is generally overcome with statistical matching of independent data sources. Typically, matching is conducted with any information common to both sources with a limited (or unknown) degree of success. In this paper Grade Correspondence Analysis (GCA) is investigated as a tool to improve the matching process. GCA draws out the relationships between the common variables to enable the sample to be partitioned into more homogeneous groups, prior to matching. An evaluation of alternative methods is conducted using datasets from the UK Family Expenditure Survey (FES), which is unusual in containing both income and expenditure at a detailed level of disaggregation. Imputed expenditures are compared with actual expenditures through the use of indirect tax simulations using the UK microsimulation model, POLIMOD. The most successful methods are then employed to enhance income data from the Family Resources Survey (FRS) and the synthetic dataset is used as a microsimulation model database.

JEL: C81 D10

KEYWORDS: Statistical matching; clustering; income and expenditure micro-data; microsimulation.

Combining household income and expenditure data in policy simulations

Holly Sutherland, Rebecca Taylor and Joanna Gomulka¹

Introduction

Statistical matching and related dataset enhancement techniques (data “fusion”) have been used when a single source of micro-data does not contain all the information necessary for a particular task. The use of such techniques may also form part of a wider strategy for improving the coherence of national data collection and for the efficient and effective use of limited data resources. However, any review of practical applications of the methods, such as Cohen (1991) reveals that no major study has consisted simply of a straightforward application of a chosen algorithm. There is inevitably a substantial degree of *ad hoc* and problem-specific treatment. Despite over 25 years of history, statistical matching remains more of a craft than a science.

Our particular problem is to create a synthetic micro-dataset containing information on both household incomes and expenditures. From the policy perspective both types of information are necessary for the analysis of the effects of the combination of direct and indirect personal taxes (Office for National Statistics, 2000; Redmond et al, 1998; Salomäki, 1996). Another specific example might be the imposition of Value Added Tax (VAT) on children’s clothes (which are currently zero-rated in the UK), combined with a compensation package for poor families which is routed through the social protection system.

Most countries do not have single sources of micro-data including high-quality disaggregated information on both incomes and expenditures. The income data in most Household Budget Surveys is usually very limited and of relatively low quality. However, the UK Family Expenditure Survey (FES) is an exception. It contains both types of information and has been used for many years as the database of official analyses of both household income and expenditures.² The FES offers an opportunity to evaluate experiments with imputation methods that may be applied to other datasets. The eventual aim is to impute expenditure variables into a second household survey dataset for the UK: the Family Resources Survey (FRS).³ The advantages of using the FRS as database for policy simulation using a tax-benefit model include a much larger sample size⁴ and detailed information necessary for the simulation of welfare benefit entitlement (such as information on savings).

Section 1 provides an overview of statistical matching principles and introduces the method to be used in this study. Common variables available in both datasets include variables

¹ The research reported in this paper was supported by the ESRC *Analysis of Large and Complex Datasets* (ALCD) Programme (H519255052). Data from the *Family Expenditure Survey* are Crown Copyright. They have been made available by the Office for National Statistics (ONS) through the Data Archive and are used by permission. Data from the Family Resources Survey have been made available by the Department of Social Security (DSS) through the Data Archive. The ONS, the DSS and the Data Archive bear no responsibility for the analysis or interpretation of the data reported here. Thanks are due to Neela Dayal and Lavinia Mitton for considerable preliminary work. This is a substantially revised version of a paper presented under the title of “Creating Order out of Chaos? Identifying Homogeneous Groups of Households across Multiple Datasets” to the 26th General Conference of the *International Association for Research in Income and Wealth* in Cracow. We are grateful for the comments received there. However, the usual disclaimers apply.

² See Office for National Statistics (1996).

³ See Department of Social Security (1997).

⁴ FRS is nearly four times the size of the FES: the 1995/6 GB samples contain 26,435 and 6,690 households respectively.

describing household characteristics and describing incomes. Although one might imagine a straightforward relationship between household income and expenditure, two factors inhibit its identification. First, there are measurement problems: incomes and expenditures are measured over different reference periods. The period for most expenditures in FES is very short: two weeks (Dayal et al, 2000). This means that while on average measured expenditures correspond to those in the population, in any one household we may observe atypical patterns such as zero expenditures on food or a very high proportion of fortnightly spending on a durable purchase (such as a car). Secondly, a linear relationship between household income and expenditure “can only be expected for a class of families homogeneous as regards tastes and needs and making their purchases on the same market” (Allen and Bowley, 1935, page 37). Our method of identification of Allen and Bowley’s ‘classes of families’ is described in section 2 of this paper.

Section 3 explains the matching process once the homogeneous groups have been identified. Section 4 presents an evaluation of alternative methods by comparing policy simulation results using both actual and imputed data from the FES. The impact of the policy changes is simulated using POLIMOD, the Microsimulation Unit’s UK tax-benefit microsimulation model (Redmond et al, 1998). Section 5 presents POLIMOD results using FRS and imputed expenditures and section 6 concludes.

1 Statistical matching of income and expenditure: the state of the art

Statistical matching started in the early 1970s. To our knowledge Okner (1972) is the earliest generally available publication on the subject. Cohen (1991) and Baker, Harris and O’ Brien (1989) provide reviews of techniques and practical applications in the field. The problem in a nutshell is as follows. We have sample A with variables (X,Y) and sample B with variables (X,Z). We want to create data C with variables (X,Y,Z) by merging records from A and B with close values of X. This is legitimate if Y and Z are related to each other only through X, i.e. if, conditionally on X, Y and Z are independent (Sims, 1972; 1974), or if relationship between Y and Z is known from other sources (for example, estimated from a different sample) and incorporated into the matching process (Paass, 1986). Under the assumption of conditional independence (which in practice can rarely be checked), a number of computational techniques for finding 'good' matches are available. Usually the samples are divided into cells by values of X and matches allowed only within cells.

The key problem for matching and imputation is classifying the samples into homogeneous groups, with the definition of similarity between households depending on the variables which are going to be imputed. In this case, these would be households with similar patterns of expenditures. The most commonly used method is defining groups as cells in a cross-tabulation of common variables. However, compromises have to be made in definition of cells between the desire to match records with very close (ideally, identical) values of X, and not creating cells with small numbers of observations. Thus in general, there is a need to identify groups that are different from 'straight-edged' cells, that is, groups other than those produced by a contingency table.

One method is simply to use regression on the whole sample. However, this involves an important limitation that does not apply to methods that explicitly identify homogeneous groups of households and then match individual records within the groups. The number of variables that may be estimated separately is limited. In particular, expenditure variables

either need to have no zero entries or the treatment of zeros needs to take account of the fact that they may be genuine zeros or may occur only as a result of the short period for recording spending. Where many variables are needed at a high level of disaggregation and where zero expenditures are important to identify from a policy point of view, individual record matching is the only suitable approach.⁵

In order to identify groups of similar households within which to match across datasets we have used Grade Correspondence Analysis (GCA). The method draws out the relationships between the common variables which are subsequently used to define clusters. Households are ordered to maximise the dependencies between household characteristics (as measured by the common variables) so that adjacent households are relatively more similar than households further apart. As far as we are aware, this method has not been used before in a statistical matching exercise, although it has been applied successfully in a number of practical studies of other types, for example Ciok et al. (1997), Ciok (1998), Szczesny and Pleszczynska (1997), Szczesny et al. (1998).

GCA as it applies to two-dimensional contingency tables is fully documented in Ciok et al. (1995). The purpose of the algorithm is to reorder the rows and columns of the table in a way which maximises a certain measure of dependence between the tabulated variables, namely the Spearman's Rho (r^*). It proceeds by alternate permutations of the rows and columns. The rule for choosing the next permutation guarantees that r^* is increased at each step. The algorithm stops when the rule cannot produce further improvements. Termination of the process does not mean that the greatest possible r^* has been reached. At the moment there is no method which would guarantee achieving the globally maximal r^* , short of trying all possible permutations of rows and columns, which is not a practical proposition for even moderately sized tables. In practice we increase the chances of finding the maximal r^* by running the algorithm a number of times on tables that have different starting orderings. Typically, before the algorithm is applied the rows and columns of the table are reordered randomly.

The primary effect of GCA is maximising r^* .⁶ Another effect is putting rows with similar values of characteristics close to each other. Thus, the relatively similar groups emerge. For more information about GCA see Taylor (2000).

2 Identifying homogeneous groups of FES households

GCA is used in three steps. First, using only variables common to both FES and FRS (except income) we attempted to define relatively homogeneous groups of households for FES95/6. Five types of common variable are considered:

1. Basic information about the sample (*sampling month* and *region*).
2. Demographic information: some variables are at the person level (*age*, *sex* and *marital status*) and some are at the household level (*household composition* - using two alternative definitions of children; *age of oldest person*; *sex of oldest person*).

⁵ For example, we might wish to explore the impact on poor families of an equalisation of tax treatments of petrol and diesel fuel. This would require that spenders and non-spenders on each type of fuel could be identified separately.

⁶ However, r^* cannot be compared across tests that use different variables.

3. Household dwelling descriptions (*category of dwelling; tenure type; total number of rooms*).
4. Potentially useful linking variables describing ownership of *durables* of various types.
5. Individual labour market activity, at the person level (*usual weekly hours, employment status; socio-economic group of the head of household*).⁷

A priori, the relative effectiveness of any combination of variables, or transformations of them, is unknown. As recommended by Baker, Harris and O'Brien (1989) to improve matching efficiency the pool of variables was chosen using regressions. The variables that exhibited predictive power across four selected expenditure categories were selected. The expenditure categories - items with VAT at 17.5%, items exempt from VAT, food, and transport – were chosen on the basis of their lack of zero values. The regressions not only eliminated the variables exhibiting the weakest relationship to expenditures, they were also informative about the form and level of aggregation of the selected variables.

Interestingly, the regression analysis eliminated variables that are regularly used in statistical matching. These included the age of the head of the household, and dummy variables counting the numbers of adults and children disaggregated by age. Some of the information was rejected in its most disaggregate form. For example, region performed better when grouped into three categorical variables by the average level of total expenditure in each region. The employment status information was most effective when the categories employed and self-employed, and sick, unemployed, and unoccupied were merged together to form two categorical variables instead of five. For car use, a continuous variable (number of cars) was rejected in favour of categorical variables to pick up a non-linear relationship between expenditure and car use (0, 1 and 2+ cars). In contrast, the continuous variable counting the number of children in the household was chosen over a categorical variable indicating the presence of any children.

Many alternative sets of cluster definitions were produced through experimentation, and to help choose between them GCA was used to identify the cluster sets that showed the strongest relationships with spending behaviour. (At this point we had no direct evidence that any of our clusters were indeed related to expenditure patterns.)

First, GCA was applied to disaggregated expenditure variables to define groupings of households according to their expenditure patterns. Expenditure variables were defined on the basis of their use by POLIMOD for indirect tax simulation.⁸ The optimal ordering of expenditure categories identified by GCA was found to be meaningful in the sense that items of expenditure at one end of the ordering were associated with high proportions of expenditure in poorer households (tobacco products, zero-rated food, domestic fuel) while categories of spending at the other end of the ordering might be expected to make up a higher proportion of total expenditure in more wealthy households (motoring expenditure, wine and champagne, household services). It is tempting to conclude that GCA can identify links between categories of expenditure that may help to distinguish between households with different lifestyles and spending patterns.

⁷ See Dayal et al (2000) for details of how these variables were constructed. Stringent criteria for the sameness of variables were used and although the two surveys are similar in many respects it proved surprisingly difficult to define identical variables in each.

⁸ The 27 variables are listed in Appendix 1 and their derivation is described in Mitton (1998). Our method would have been equally appropriate if we had wished to impute directly from the FES database the full set of 400 disaggregated expenditure variables.

Finally, GCA was used to see how the two types of groupings relate to each other. The sets of household characteristic clusters that exhibited the strongest relations with the expenditure groupings were chosen for subsequent matching attempts. For more information and illustrations of how this was done see Taylor et al (2001) and Taylor (2000). Appendix 2 shows the three chosen cluster definitions and the sample proportions in the FES and FRS datasets.

3 Matching

Having identified similar groups of households in each of the datasets the households within each group are ranked by household income. Then they are sequentially matched, starting with the lowest income household from a donor cluster matched to the lowest income household in the equivalently defined cluster in the recipient file. We call this ‘rank by income’ matching. Households from any recipient cluster receive expenditure variables only from the corresponding donor cluster. If the size of the donor and recipient clusters were the same, then one donor household would provide expenditure variables for one recipient household. In any other case, when the sample size of the clusters in the donor and recipient files are unequal, the appropriate proportions of variables from neighbouring households are taken from donor to recipient dataset.

The ‘rank by income’ method can be augmented using common variables that have not been used to define clusters. Within clusters, households can be ranked by another variable (e.g. the number of adults in the household) before ranking by income within the new sub-categories. Matching is performed across the re-ordered clusters. However, re-ranking the households may disrupt the matching. Since the number of households in a corresponding sub-category in each dataset is unlikely to be identical, this will allow high-income households to be matched with low-income households. A second method avoids this problem by splitting clusters by the extra information and then ranking by income and matching proceeds within the split clusters. The potential disadvantage of splitting the clusters is the decrease in the donor sample size.

4 Evaluation using Family Expenditure Survey

No formal statistical tests exist to distinguish between the alternative sets of imputed data. Our evaluation relies on comparison of microsimulation model results when imputed expenditure data are used, in relation to results based on actual expenditure data. Comparisons were made using FES data from a two-year sample (94/5 and 95/6). Expenditures were imputed from one (random) half of the combined sample into the other half. Then the roles of donor and recipient were reversed, providing an actual and an imputed set of expenditures for each household in the combined sample. Many alternative sets of imputed data were generated.

The main method of evaluation was to compare the outcome of simulated changes in indirect tax policy across the distribution of equivalised household income (using the modified OECD equivalence scale and counting each household once in the ranking). Although the expenditure data varies between each comparison, a common micro-database for household income and characteristics variables is used in all of the estimations. FES data are updated to

2000/1 prices and incomes, and use 2000/1 UK tax and social security policy as a starting point. All simulations use the data re-weighted to represent the national population. In calculating the impact of the tax change, it is assumed that households do not change quantity of goods bought.

The first policy change involved simulating the removal of Value Added Tax (VAT) from all goods (“*novat*”). This was designed to test whether the imputed distribution of goods and services that attract VAT reproduces the distribution of actual expenditure on this group of goods and services. The second simulation - called “*revneut*” - found the revenue-neutral rate of VAT that would apply if all goods and services were subject to the same, uniform rate (many goods are currently zero-rated or exempt; most VAT is charged at 17.5%). This resulted in some households gaining and some losing, depending on their spending patterns. The proportion in each group, as well as the estimated revenue-neutral VAT rate provided us with examples of estimates that are commonly used by policy makers, and which at the same time were likely to be sensitive to differences in patterns of expenditure.

The very extensive comparisons of POLIMOD output using actual FES expenditures and different imputed datasets are described in Taylor et al (2001). Box 1 lists the main imputations discussed here. These variations were designed to test whether:

- (a) pre-match clustering of households improves on un-clustered ‘rank by income’ matching and there were any differences in the three alternative sets of cluster definitions.
- (b) the choice of variable used to rank within clusters has any affect
- (c) the performance of pre-match clustering can be improved by introducing additional information within each cluster.

Box 1: Summary of imputations		
Name	Clusters	Ranking and splitting within clusters
Impf1	Unclustered	Rank by household disposable income (hhdispy)
Impf2a	Test42b.1	Rank by hhdispy
Impf2b	Test42b.2	Rank by hhdispy
Impf2c	Test41.1	Rank by hhdispy
Imp2aex	Test42b.1	Rank donor by total expenditure, receiver by hhdispy
Impf7	Test42b.1	Rank by children then hhdispy
Impf9	Test42b.1	Customised split then rank by hhdispy
Impf13	Test42b.1	Split by children then rank by hhdispy

Sampling error is one influence explaining differential effects across datasets. The size of this effect was used as a benchmark to provide some indication of the importance of the differences between imputations and between results based on imputed and actual data. The actual data were split randomly into two equally sized samples and used to generate two equivalent sets of POLIMOD results.

Aggregate results are summarised in Tables 1 and 2 and selected results are illustrated in Figures 1 to 5.

(a) Pre-match clustering and differences between cluster definitions

Rows 4, 5 and 6 of tables 1 and 2 show results for three pre-match clustered imputed datasets. The cluster definitions are described in Appendix 2. These results may be contrasted with those based on actual expenditure data (row 1) and those based on expenditures imputed by ranking by income and matching across whole datasets, without clustering (row 2). We can see that:

- Mean total expenditure is very similar across all five variants (to be expected, given the imputation method), although the variation is somewhat less in the clustered datasets than in the un-clustered or the actual expenditure data (Table 1).
- The share of total expenditure across the distribution of household income is much closer in the three clustered datasets to the actual than in the un-clustered dataset, which over-estimates the share of the high income groups and underestimates the share of expenditure among those on low incomes (Table 1).
- The aggregate effects of simulated policy changes are also similar across the five alternatives. Estimates of the change in indirect tax revenue when VAT is abolished, the revenue-neutral rate of uniform VAT and the proportion of households losing on the introduction of a uniform rate are all within 1.5% of the actual estimates (Table 2).
- There is little to choose between the three sets of clusters in terms of the aggregate results shown in these tables.

The estimated impact of *novat* across the income distribution was analysed using boxplots of the VAT paid by households disaggregated by household income decile groups. Figure 1 shows the distributional impact using the five alternative expenditure datasets. The line through the centre of the boxes in the plots represents the median, the box represents the inter-quartile range, and the whiskers reach out to the lowest and highest values (excluding outliers which are defined as more than 1.5 times the third quartile above the box and 1.5 times the first quartile below the box).

- There is a noticeable difference in the median of un-clustered estimates (VAT1) and the actual values in the bottom and top decile groups.
- There is little to choose between the clustered estimates. Although some perform better than others in some decile groups, performance is not consistent across the whole income distribution. On the whole, Impf2a (VAT2A) looks closest to the estimates based on actual data.

While the imputations may perform reasonably well for the whole sample, the same is not always the case for sub-groups. Figure 2 shows the median change in VAT within each decile group under *revneut*, focusing on households with children and shows that:

- None of the imputations capture the high median loss among households with children in the top decile group.
- The un-clustered data (VAT1) do not replicate the actual data well.
- The different cluster definitions result in different estimates.
- Impf2a (VAT2A) performed less well than the other two imputed datasets based on pre-match clustering.
- The best-performing dataset is Impf2c which was based on clusters that include presence of children in their definition.

(b) Choice of ranking variable

The baseline definition of income used as the ranking variable before matching is household disposable income after committed expenditures i.e. minus income tax, National Insurance

contributions and housing costs.⁹ Four alternative income definitions are used to explore the sensitivity of the estimates: disposable income before housing costs, disposable income after housing costs per adult and equivalised disposable income after housing costs. Two alternative equivalence scales were used: the modified OECD scale and McClements scale. We found that the particular definition of income does not affect the distributional estimates significantly. However the conclusions are quite different when two differently defined variables are used to rank the donor and receiver samples. Figure 3 compares the effect of un-clustered matching (VAT1) and with clustered matching using the same income variable to rank both samples (VAT2A) with an imputation using the same clusters but with households ranked within clusters by income in the recipient dataset and total expenditure in the donor dataset. This clearly demonstrates that in nearly all respects the imputed data using different ranking variables (VAT2aex) is less like the actual data than either of the alternatives. The results suggest that caution should be exercised in ranking the donor and receiver samples by different variables, but that if the income variable used for ranking is the same in both datasets, its exact definition is of little importance.

(c) Adding information to pre-match clustering

These experiments are illustrated with reference to the sub-sample of households with children. Three distinct ways of introducing information on children are explored. The first is to include the presence of children in the cluster definition (Impf2c). The second is to rank households by the presence of children within each cluster, then rank by income within the categories in each cluster and then match within whole clusters. (In this case the base cluster definition - Impf2a - does not include children.) The third method is to split clusters (Impf2a) according to the presence of children and to match within sub-clusters. Figure 4 contrasts results for *novat* using the three approaches, with the results using actual data and results using pre-match clustering not involving information about children.

- Ranking by children and then matching (VAT7) performs very badly. This is because this method allows some high income households with (without) children to be matched with low-income households with (without) children.¹⁰
- Clustering using child information (VAT2C) is better than clustering without child information (VAT2A) when households with children are the focus.
- Splitting clusters (that do not depend on children) by presence of children (VAT13) performs best for this sub-group.

To consider the results in relation to sampling error, the sample is split randomly into two. The actual and imputed data estimates calculated from sample 1 are compared directly to the actual data estimates from sample 2. The change in indirect tax from policy *revneut* is examined over quantiles of the distribution of household income. The households that gain and lose because of the policy change have been separated, and Figure 5 only shows the losers. Kernel regression is used to smooth the curves to reduce the impact of extreme values and to make the charts readable.¹¹ All variables have been weighted to represent the population. The figure shows that there is a substantial amount of variation between the two samples of actual data compared with the differences between results using the actual and imputed estimates. There is little difference between the sets of results using imputed data in Figure 5.

⁹ For a full definition see Dayal et al (2000).

¹⁰ For an explanation of how this occurs see Taylor et al (2000).

¹¹ A linear weight of decreasing importance is used to include the impact of the surrounding points.

Three alternative imputation methods were selected to be implemented with FRS data: Impf2a (pre-match clusters not including children in their definitions), Impf2c (pre-match clusters including child information) and Impf13 (the same as Impf2a but with clusters split into households with and without children).

5 Imputation of expenditures into Family Resources Survey data

Three policy simulations are used to provide results using FRS with imputed expenditures that could be compared with POLIMOD results using actual FES data. As well as the two already considered (*novat* and *revneut*), a third policy change - the imposition of VAT on children's clothes (*child-cloth*) - is used. This is designed to test the performance of the imputations in relation to a relatively small component of expenditure of particular relevance to an identifiable sub-group of households.

In comparing results using FES with those using FRS and imputed expenditures it is important to remember that the expenditure data are not the only source of difference. We have found that, although the surveys are similar in many respects, the distributions of household income are significantly different (Dayal et al, 2000). FRS incomes are lower than FES incomes on average, but for some sub-groups the opposite is the case.¹² We would therefore expect the composition of income decile groups in the two datasets not to be identical. For this reason we would not expect the expenditures of (say) the bottom income decile group in the FRS sample to be the same as the expenditures in the bottom decile group of the FES. In addition, routine data adjustments (e.g. re-weighting to correct for differential non-response) and the process of policy simulation (e.g. the modelling of non-take-up of some social security benefits) may either exacerbate or mitigate these underlying differences. Thus, in the comparisons that follow we should not expect to find the same results using the alternative data sources. Rather, we can conclude that the imputations are sufficiently robust for them to be considered adequate for the particular purpose if the two sets of results lead us to the same policy conclusions.

Table 3 shows the aggregate results using FES and enhanced FRS data. The results from the *novat* policy change suggest that the (weighted) imputed datasets contain less expenditure that attracts VAT than the (weighted) FES data. This is consistent with the somewhat lower revenue-neutral uniform VAT rates obtained when using imputed data and a slightly lower proportion of households losing when the revenue-neutral rate is implemented. However, all these differences are small and there is little to choose between the imputations. Figure 6 shows the distribution of VAT across income decile groups, which confirms the similarity of the datasets at this general level. Figure 7 shows the average gain or loss in each decile group from a revenue-neutral uniform VAT rate. The imputed datasets all pick up the general pattern of larger gains at the top of the income distribution, but tend to underestimate both the loss at the bottom and the gain at the top. Impf2c seems to follow the actual relationship most closely.

The differences between the results for the simulation of VAT on children's clothes shown in Table 3 are larger. The cost estimate is between 8% and 15% larger for the imputed datasets compared with the real data. The imputed data also show a larger proportion of households being affected than the actual data. Although the FES data show that 15% of spending on

¹² See also Frosztega et al (2000).

children's clothes occurs in households without children, the estimates using imputed data suggest that this proportion is even higher. This is particularly the case for Impf2a which at no stage controls for the presence of children in the matching process. Without this control we find that 56% of the change falls on households without children and that only 42% of households with children are affected (compared with 58% in the real data). When the presence of children is controlled for, the results using imputations are closer to the actuals. However, the results shown here do not point to which of the two methods that control for children is to be preferred. On the one hand, Impf2c closely matches the actual proportion of households with children affected (59% compared with 58%), while the Impf13 estimate is too high (65%). On the other hand, the share of the impact on households without children is only a little too high for Impf13 (16% compared with 15%) but nearly double what it should be using Impf2c (29%).

Figure 8 shows the average VAT paid on children's clothes across the deciles of the household income distribution. The actual relationship is quite flat, but all the imputations fail to capture fully the flattening at the top of the distribution. Impf2a, which does not control for children behaves particularly badly in this respect. This is even clearer in Figure 9, which plots the same information for households with children only. Impf2a underestimates the effect, and here it is clear that it does so regardless of income level. The other two imputations follow the actual more closely, but at both top and bottom of the distribution Impf13 overestimates, and Impf2c underestimates the impact. Overall, the policy analyst would probably draw similar conclusions from simulations using FRS data and Impf13 or Impf2c to those using FES data.

Clearly, if changes affecting a specific group of households are to be modelled, the imputation method must take account of the characteristics of that group. In the case of tax on children's clothes, children must be explicitly controlled for in the imputation. However, there are many groups of interest from a policy point of view and it is unlikely that a general imputation method could be found that simultaneously anticipated all such groups and was able to incorporate sufficient information. (The methods used here involved finding a balance between the number of variables and categories taken into account in the cluster definitions and splits, and in the number of observations in each group to be matched. If the samples are divided into too many groups before matching, it is made more likely that relatively high-income households may be matched with relatively low-income households.)

In addition, a full analysis of VAT on children's clothes might require us to look at sub-groups - say lone parent families, or those with pre-school children. In such cases it is most unlikely that imputation methods that simply controlled for children as a whole would provide results that were close to those from the original data.

Furthermore, the example we have chosen - spending on children's clothes by households with children - might be expected to behave better than most. The households that spend on children's clothes are mainly identified by the presence of children - a characteristic that is observed and can be controlled in the matching process. In the case of spending on many other goods - for example, cigars or diesel fuel, each of which are taxed in specific ways that may be of policy interest - it is more difficult to find observed characteristics that predict expenditure.¹³

¹³ An indication of the categories of spending that that may be difficult to predict using our method is provided in Appendix 3 which compares the decile shares of each category of imputed expenditure for FRS (using equivalised income deciles) with corresponding decile shares from FES data.

6 Conclusions

It is clear that spending patterns, as captured by FES data, vary considerably from sample to sample. This can be explained by the short reference period and the dominating influence of atypical expenditures. It means that it is particularly difficult to predict expenditure patterns that capture micro-level diversity, as exhibited in the actual data. Imputed expenditure variables are only ever an adequate second-best substitute for actual data when the variables are used at a sufficient level of aggregation to mask differences that are not controlled by the imputation procedure. Since we cannot say *a priori* what this level of aggregation should be, we can only be confident in the imputations when the dimensions that are important to the analysis have been controlled for.

Our results comparing pre-clustered and un-clustered rank-by-income matching suggest that the identification of homogeneous groups using GCA is a fruitful approach, but that the dimensions of importance to the subsequent policy analysis should be included in the cluster definitions.

If the purpose is only ever to produce results at an aggregate level or for large groups identified by variables common to both datasets (and used in the matching process) then our method seems adequate in relation to the benchmark provided by the results using two samples from the same dataset. However, in many practical 'expenditure-to-income' imputation exercises, the donor dataset does not contain precise or good quality income variables. Our experiments suggest that it is important for the ranking variable to be the same in both donor and recipient dataset. Where this is not possible it is likely that the imputations are of lower quality than those produced in our experiments.

References

- Baker K, P Harris and J O'Brien (1989), 'Data fusion: an appraisal and experimental evaluation', *Journal of the market research society*, 31(2), pp153-212.
- Central Statistical Office (1995), *Family Spending. A report on the 1994-95 Family Expenditure Survey*, London: HMSO.
- Ciok A (1998), 'A comparative study of exploratory methods applied to car switching data', *Bulletin of the Polish Academy of Sciences – Technical Sciences*, 46, pp 133-148.
- Ciok A, W Bulhak, R Skoczylas (1997), 'Exploration of control-experimental data by means of grade correspondence analysis', *Biocybernetics and Biomedical Engineering*, 17, pp 101-113.
- Ciok A, T Kowalczyk, E Pleszczyńska and W Szczesny (1995), 'Algorithms of grade correspondence – cluster analysis', *Archiwum Informatyki Teoretycznej i Stosowanej* (The Collected Papers on Theoretical and Applied Computer Science), 7, pp 5-22.
- Cohen M L (1991), 'Statistical Matching and Microsimulation Models', in C F Citro and E A Hanushek (eds.), *Improving Information for Social Policy Decisions. The Uses of Microsimulation Modeling* vol. II, National Academy Press, Washington D.C.

- Dayal N, J Gomulka, L Mitton, H Sutherland and R Taylor (2000), 'Enhancing Family Resources Survey income data with expenditure data from the Family Expenditure Survey: data comparisons', Microsimulation Unit Research Note, MU/RN/40, Department of Applied Economics, University of Cambridge.
- Department of Social Security (1997), *Family Resources Survey: Great Britain 1995-96*, London: The Stationary Office
- Frosztega M and the Households Below Average Income team (2000), *Comparisons of income data between the Family Expenditure Survey and the Family Resources Survey*, GSS Methodology Series no. 18, Office for National Statistics.
- Mitton L (1998), 'POLIMOD: the calculation of VAT and Excise duties on household expenditure', Microsimulation Unit Research Note MU/RN/29, Department of Applied Economics, University of Cambridge.
- Okner B (1972), 'Constructing a new microdata base from existing microdata sets: The 1966 merge file', *Annals of Economic and Social Measurement*, 1, pp 325-362.
- Office for National Statistics (1996), *Family Spending: A report on the 1995-96 Family Expenditure Survey*, London: The Stationary Office.
- Office for National Statistics (2000), 'The effects of taxes and benefits in household income 1997-98', *Economic Trends*, (April), London: The Stationary Office.
- Paass G (1986), 'Statistical Match: Evaluation of Existing Procedures and Improvement by Using Additional Information', in G H Orcutt, J Merz and H Quinke (eds.), *Microanalytic Simulation Models to Support Social and Financial Policy*, Elsevier Science Publishers B.V.
- Redmond G, H Sutherland and M Wilson (1998), *The arithmetic of tax and social security reform: a user's guide to microsimulation methods and analysis*, Cambridge University Press.
- Salomäki A (1996), 'Including Consumption Expenditure and Welfare Services in a Microsimulation Model', in A Harding (ed.), *Microsimulation and Public Policy*, Elsevier, Amsterdam.
- Sims C A (1972), 'Comments and rejoinder', *Annals of Economic and Social Measurement*, 1, pp 343-345, 355-357.
- Sims C A (1974), 'Comment', *Annals of Economic and Social Measurement*, 3, pp 395-397.
- Szczesny W, A Ciok and E Pleszczyńska (1998), 'Clustering land districts according to their farm magnitude repartition', *Statistics in Transition*, 3, 757-768.
- Szczesny W and E Pleszczyńska (1997), 'A grade statistics approach to exploratory analysis of the HSV data', *Biocybernetics and Biomedical Engineering*, 17, pp 235-245
- Taylor R (2000), 'Guidelines for Identifying Clusters Using Grade Correspondence Analysis: Practical and Technical Issues', Microsimulation Unit Research Note MU/RN/39, Department of Applied Economics, University of Cambridge.
- Taylor R, H Sutherland and J Gomulka (2001), 'Using POLIMOD to evaluate alternative methods of expenditure imputation', Microsimulation Unit Research Note MU/RN/38, Department of Applied Economics, University of Cambridge.

Table 1: Share of total household expenditure by equivalised income decile group and imputation method

Imputation	Mean	S.D.	Equivalised household income decile groups									
			1	2	3	4	5	6	7	8	9	10
1 ACTR	13811.71	12743	4.80	3.92	5.60	7.46	8.69	9.94	11.66	13.56	14.62	19.73
2 impf1	13811.65	12225	3.34	3.57	4.94	6.66	7.99	9.61	11.07	13.25	15.63	23.94
3 impf2a	13803.41	11521	4.56	3.96	5.60	7.22	8.50	9.97	11.58	13.38	15.20	20.04
4 impf2b	13801.62	11622	4.56	4.00	5.48	7.23	8.54	10.02	11.66	13.28	15.33	19.90
5 impf2c	13788.75	11486	4.58	3.92	5.49	7.33	8.82	10.05	11.29	13.53	14.85	20.15
6 impf2aex	13803.41	12675	2.74	3.24	4.51	6.08	7.34	8.99	10.77	13.03	16.03	27.27
7 impf7	13803.41	11428	5.22	4.47	6.18	8.22	9.18	10.73	11.61	13.16	14.29	16.94
8 impf9	13807.94	11643	4.56	3.95	5.65	7.26	8.73	10.07	11.38	13.43	15.02	19.95
9 impf13	13804.60	11542	4.61	4.00	5.51	7.41	8.73	9.96	11.55	13.56	15.04	19.64

Notes: Amounts are £/year at 1994-6 prices. The modified OECD equivalence scale has been used to rank households.

Table 2: POLIMOD output by imputation method, using FES

Imputation	policy change:	“novat”	“revneut”	
		change in indirect tax £ million/year	uniform rate of VAT %	households losing %
1 ACTR		37253	9.98	51.40
2 impf1		37479	9.96	52.17
3 impf2a		37393	10.04	51.22
4 impf2b		37432	10.04	51.28
5 impf2c		37354	10.06	50.72
6 impf2aex		37429	10.05	52.93
7 impf7		37303	10.07	51.18
8 impf9		37413	10.05	51.30
9 impf13		37398	10.08	51.42

Notes: Results are expressed in terms of 2000/1 prices and incomes and are weighted to represent the national population.

Table 3: POLIMOD output for three policy changes, using FES and enhanced FRS

Data / imputation	Policy change:	“novat”	“revneut”		change in VAT £ million/year	“child-cloth”		
		change in indirect tax £ million/year	uniform rate of VAT %	households losing %		% affected		% falling on households without children
						all	with children	
FES 94/5+ 95/6		37253	9.98	51.4	625	22.6	58.2	15.1
FRS + Impf2a		37021	9.85	49.9	720	28.7	41.7	55.7
FRS + Impf2c		37009	9.85	50.0	694	27.1	58.5	29.1
FRS + Impf13		37042	9.87	49.9	678	26.4	65.4	16.4

Notes: Results are expressed in terms of 2000/1 prices and incomes and are weighted to represent the national population.

Figure 1: Estimate of VAT difference for households for policy change *novat* to examine the effect of clustering the data (full sample)




Figure 2: Estimate of median VAT change for households for policy change *revneut* to examine the effect of clustering the data (households with children)

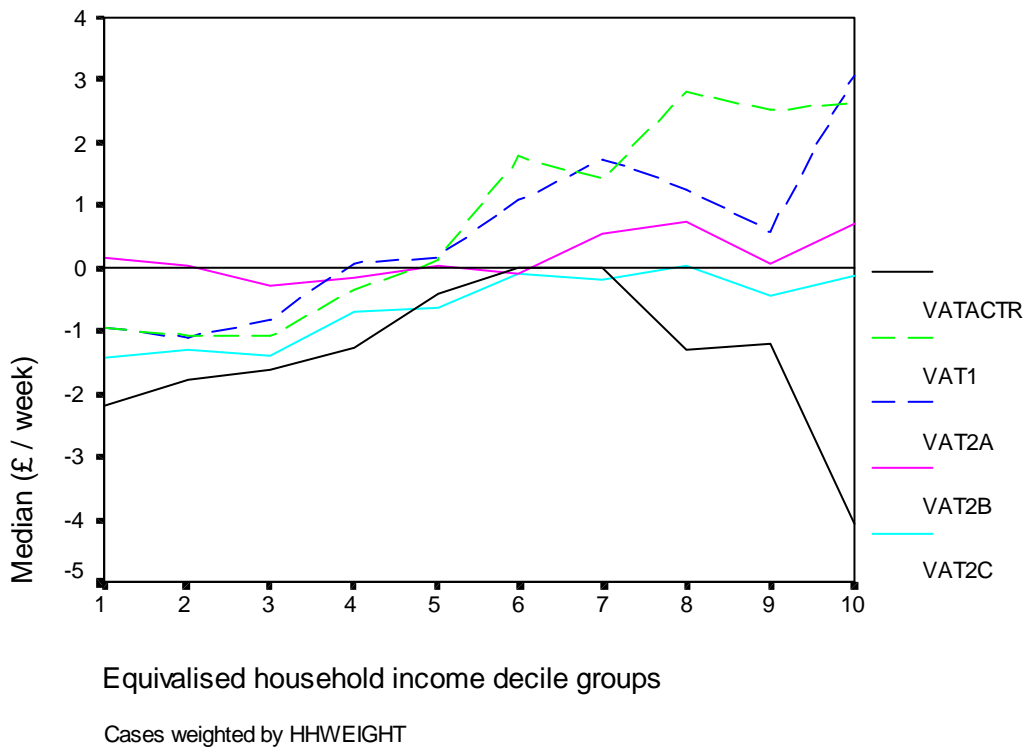


Figure 3: Estimate of VAT difference for households for policy change *novat* to examine the effect of ranking the donor sample by expenditure (full sample)




Figure 4: Estimate of VAT difference for households for policy change *novat* to compare the way that information on children is incorporated (households with children)

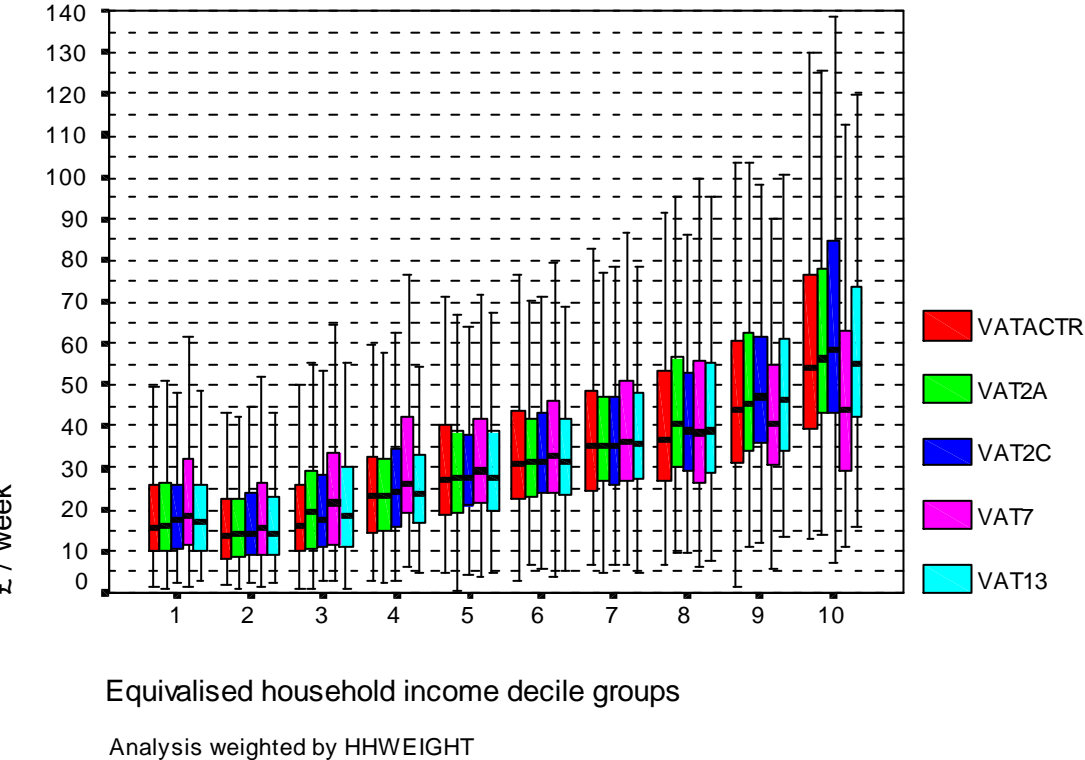


Figure 5:

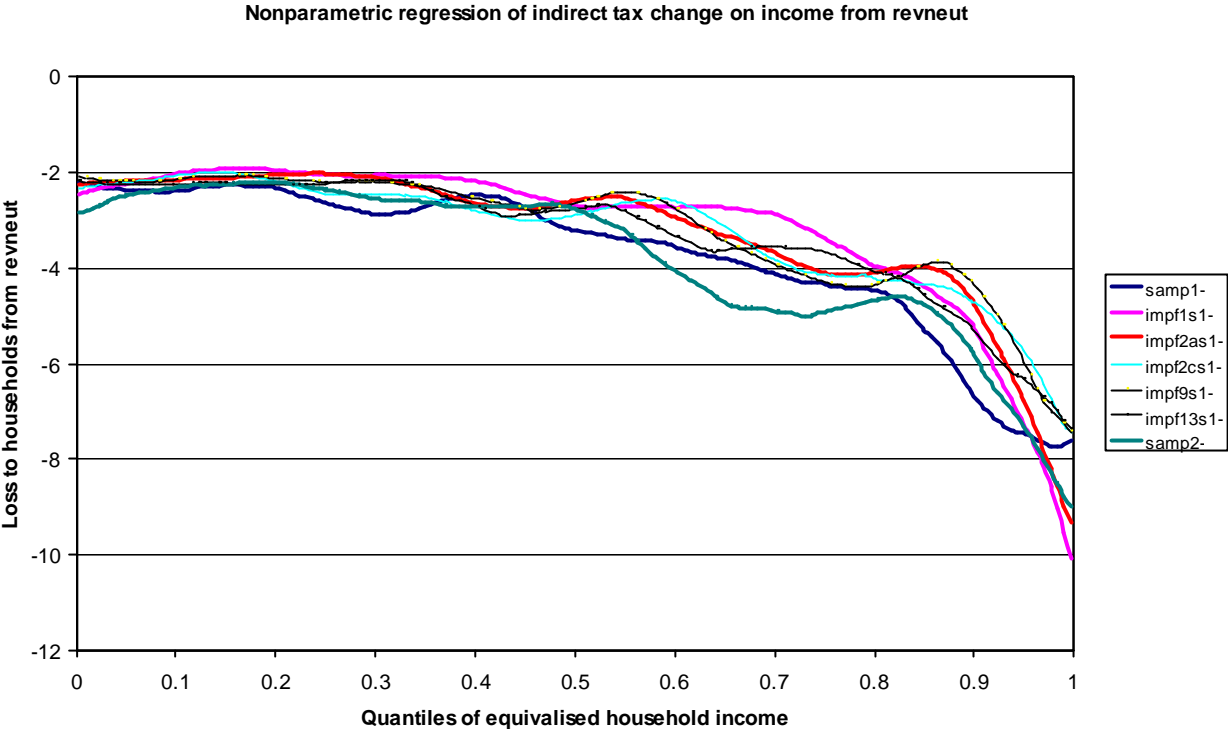


Figure 6: Mean VAT by household income decile: comparing FES actual with FRS imputed




Figure 7: Gains and losses from a uniform rate of VAT: comparing FES actual with FRS imputed

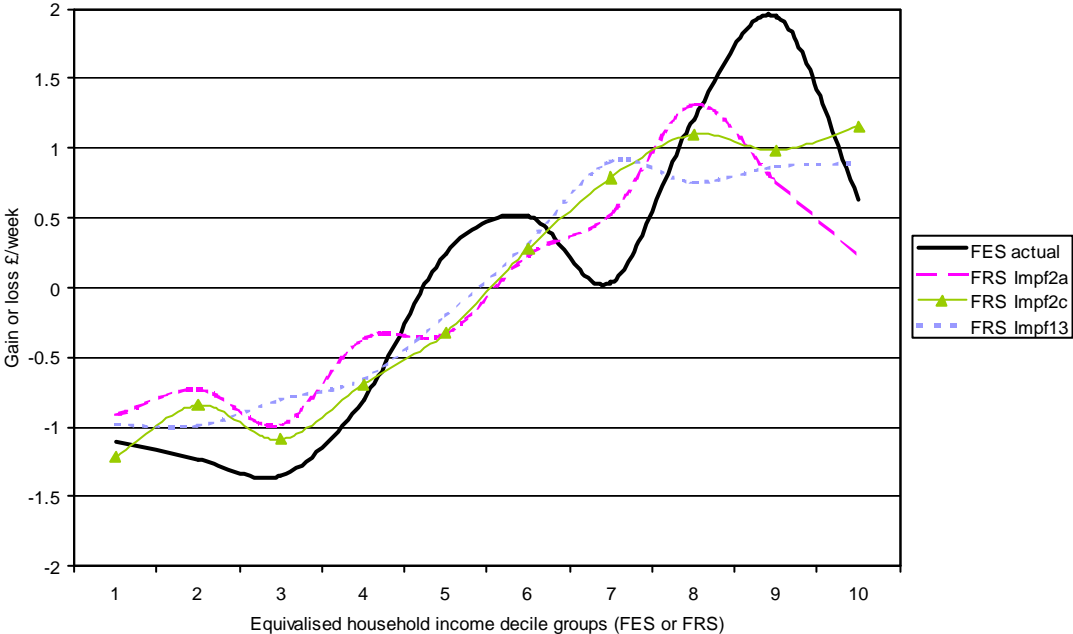


Figure 8: VAT on children's clothes: comparing FES actual with FRS imputed (all households)


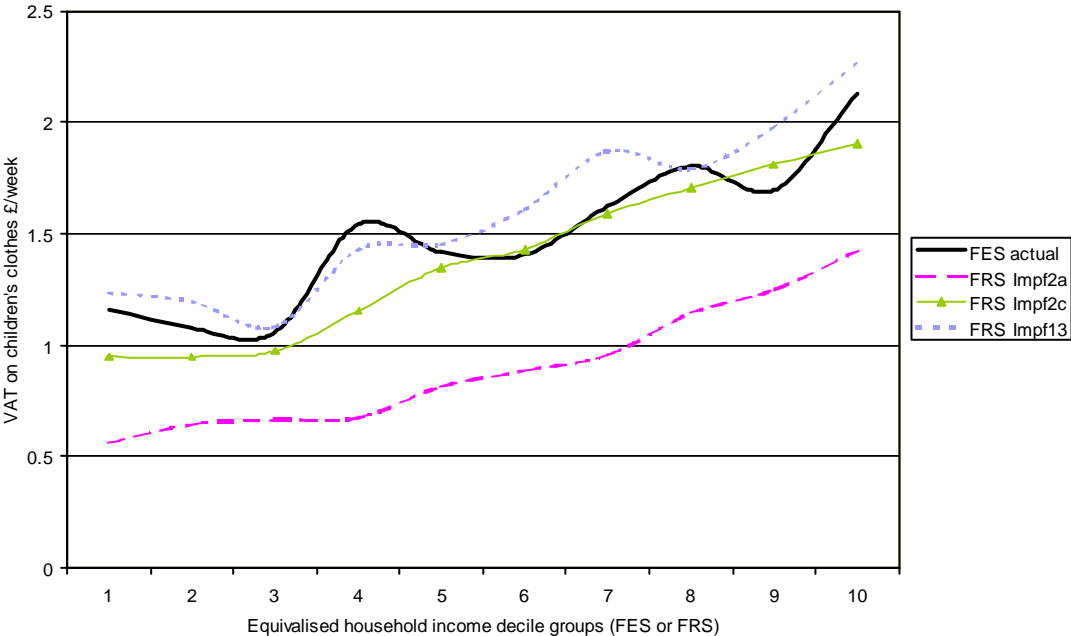


Figure 9: VAT on children's clothes: comparing FES actual with FRS imputed (households with children)



Appendix 1: Expenditure variables using FES 1995/6 by tax treatment

Expenditure category	Mean £/week	SD	% of hholds with +ve expenditure	Current tax treatment (2000/1)
1 Housing expenditure + household services + other household expenditure	23.90	54.23	95.0	VAT (17.5%)
2 Motoring expenditure	12.58	65.47	58.9	VAT (17.5%)
3 Food (which attracts VAT)	14.45	16.89	94.3	VAT (17.5%)
4 Leisure goods and services	17.82	39.07	86.2	VAT (17.5%)
5 Adult clothing and footwear	13.89	26.84	62.5	VAT (17.5%)
6 Household goods + personal goods and services	32.99	48.70	98.3	VAT (17.5%)
7 VAT-exempt goods*	45.76	113.19	98.5	No VAT
8 Food (zero-rated for VAT)	39.19	24.08	99.8	No VAT
9 Books and newspapers	4.09	5.23	91.8	No VAT
10 Domestic fuel and power	12.88	7.78	97.7	reduced rate VAT (5%)
11 Other zero-rated goods (includes transport and drugs and medicines)	8.44	85.19	79.1	No VAT
12 Children's clothes	3.55	10.36	26.6	No VAT
13 Insurance premia	10.31	10.77	87.2	insurance premium tax only
14 Beer	6.96	13.01	57.9	VAT + excise duty
15 Cider	0.30	1.57	20.5	VAT + excise duty
16 Fortified wine	0.37	1.30	19.2	VAT + excise duty
17 Wine	2.11	5.56	40.8	VAT + excise duty
18 Champagne	0.11	1.21	15.1	VAT + excise duty
19 Spirits	1.63	4.84	30.3	VAT + excise duty
20 Cigarettes	5.23	10.04	35.0	VAT + excise duty
21 Cigars	0.14	1.36	2.2	VAT + excise duty
22 Pipe tobacco	0.40	1.89	6.6	VAT + excise duty
23 Motor fuel	9.90	12.41	60.7	VAT + excise duty
24 Motor fuel (diesel)	0.79	4.25	5.2	VAT + excise duty
25 Pools stakes	0.38	1.29	18.7	excise duty
26 Other betting stakes	1.18	4.67	27.2	excise duty
27 Lottery stakes	2.30	3.18	69.7	excise duty

* includes postal services, life insurance, financial services, education, health, burial and cremation and

trade union and professional subscriptions.

Appendix 2: Cluster definitions

No	Clusters = 42b.1	%		Clusters = 42b.2	%	Clusters = 41.1	%	
		FES	FRS				FES	FRS
1	Single retired, rentunf1, no car	5.2	6.0	Single retired, no car, rentunf1	5.2	2+ cars, ownsome/ownall, workers, no children	6.2	5.8
2	Single retired, ownall, no car	4.2	4.7	Single retired, no car, ownall	4.2	2+ cars, ownsome/ownall, workers, children	6.0	5.9
3	Single other, rentunf1, no car	4.2	4.5	Single other, no car, rentunf1	4.2	1 car, ownsome, workers, no children, high	3.6	3.4
4	>1 retired, retired&other, rentunf1, no car	1.8	1.8	Single retired, no car, rent or multiples of (1)	2.0	1 car, ownsome, workers, children, high	2.4	2.3
5	1/2 retired rent, 2 retired ownall, no car	2.1	2.5	1/2 retired or 1 other, no car, rentunf1/ownall	1.9	1 car, ownsome, workers, no children, mid/low	5.1	4.6
6	mixed bag- 6	2.0	2.0	Retired/other, no/1 car, rent/ownall	2.0	2+ cars, ownsome, work&other/retired	4.0	3.5
7	Other, rentunf1/rent, no car	3.0	3.4	1/2 other, no car, rent	2.8	1 car, ownsome, workers, children, mid/low	4.8	4.2
8	1 work, 1 work&1/2 retired, rentunf1, no car	1.5	1.7	At least one worker, no car	3.8	1 car, ownall, workers	2.8	2.9
9	Single retired, ownall, 1 car	2.5	2.5	Single retired, 1 car, ownall	2.5	2+ cars, (rent, work or work&other/ret) or (ownall, work&other/ret)	3.4	3.3
10	2 retired, ownall, 1 car	4.4	4.6	2 retired, 1 car, ownall	4.4	1 car, ownsome, work&1other	4.3	4.0
11	1 work&1 retired, 1 work&1 other, no car	2.2	2.0	Mixed bag, no/1 car	2.0	2+ cars, no workers	1.3	1.3
12	Retired and/or other, not ownsome, 1 car	4.2	4.6	Other, no/1 car	2.8	1 car, (rent, workers) or (ownsome, work&other/retired), children	1.9	1.7
13	Retired and/or other, no car	2.1	2.1	at least one worker, no car, rentunf1/ownall	1.6	1 car, ownsome, no workers	2.0	2.1
14	Single worker, not rentunf1, no car	2.6	2.7	2 other, rentunf1 or 1 retired&1 other, ownall, 1 car	1.4	1 car, rent, workers, no children	2.4	2.2
15	>1 worker, no car	2.2	1.9	Mixed bag- 15	2.7	no car, ownsome, workers or work&other/retired	3.3	3.4
16	1 work&1 other, not ownsome, 1 car	2.6	2.3	1 work&1 retired, ownall or 1 work&1 other, rentunf1, 1 car	1.4	1 car, rent/ownall, work&other/retired, children	1.9	1.6
17	Workers, and retired and/or other, 1+ cars	4.2	4.2	Single worker, no car&ownsome or 1 car&rentunf1	2.0	1 car, rent/ownall, no workers, children	1.4	1.6
18	Retired and/or other, 1+ cars	3.1	3.2	1 work&1 other, ownall/rentunf2, or 1 retired&1 other, ownsome, 1 car	1.4	1 car, work&other/retired, no children	4.1	4.2
19	Workers, rentunf1/ownall, 1 car	4.5	4.6	Mixed bag- 19	2.4	no car, rent/ownall, workers or work&other/ret, children	2.0	2.0
20	Workers, rent, 1 car	2.4	2.0	2 work, no car, ownsome or 2 retired, 2+ cars, ownall	1.6	no car, rent/ownall, workers or work&other/retired, no children	3.8	3.7
21	1 work&1 other, ownsome, 1 car	3.6	3.4	1/2 workers, ownall/rentunf1/2, 1 car	2.5	no car, no workers, children	4.2	4.7
22	mixed bag- 22	2.5	2.6	1 work, rent/free or 2 work, ownall, 1 car	2.1	(no car, ownsome, no work) or (1 car, rent/ownall, other), no children	2.4	2.5
23	Workers, not ownsome, 2+ cars, and 3+ workers, ownsome, 1 car	3.3	3.1	Mixed bag, 1+ cars- 23	2.2	1 car, ownall, retired or retired&other, no children	3.9	4.0
24	mixed bag- 24	1.9	1.6	Mixed bag, 1+ cars- 24	1.7	1 car, rent, retired or retired&other, no children	1.7	1.8
25	2 workers, ownsome, 1 car	10.5	9.5	1 work&1 other, ownsome, 1 car	3.6	1 car, ownall, retired, no children (1f low/mid, 2f mid, 2 m/flow/mid)	3.7	3.9
26	Single other, 1 other&1 work, ownsome, 2+ cars	2.3	2.0	Mixed bag, 2+ cars- 26	2.6	no car, rent/ownall, other, no children	3.8	4.1
27	Single worker, ownsome, 1 car	4.8	4.3	Workers, 1 car&ownsome, or 2+ cars&ownall	2.1	no car, rent, retired or retired&other, no children (not 29/30)	2.8	3.2
28	3+ workers, ownsome, 2+ cars	1.9	1.6	All work or work and 1 other/retired, mainly ownsome, 2+ cars	2.6	no car, ownall, retired or retired&other, no children (not 31)	3.5	4.2
29	1/2 workers, ownsome, 2+ cars	8.8	8.6	2 workers, ownsome, 1 car	10.5	no car, rent, retired, no children (1 m low/mid, 2 f mid, 2 f/m, low/mid)	1.7	1.9
30				1 work&1 other, ownsome, 2+ cars	2.3	no car, rent, single retired female, no children, low/mid	3.1	3.4
31				Single worker, ownsome, 1 car	4.8	no car, ownall, single retired female, no children, low/mid	2.3	2.5
32				3 workers, ownsome, 2+ cars	1.5			
33				2 workers, ownsome, 2+ cars	8.6			

Notes: percentages shown are for combined 94/5 and 95/6 FES datasets and for 95/6 FRS. Clusters using 42b.2 were not calculated for FRS.

Variable definitions

Variable	Definition
car0, car1, car2	Dummy variables - number of cars used in HH (0, 1 or 2+ respectively)
children	Dummy variable - 1 if HH has children between the ages of 0 and 16 (or 18 if in secondary education)
dssad	Number of adults in HH
high, mid, low	Dummy variables - region of HH (high, medium or low expenditure region, as measured by CSO (1995))
ownsme, rent	Dummy variables - housing tenure, aggregated (have mortgage or partown, rent). Omitted as base = own outright/rent free
ownsme, ownall, rfree, rentf, rentunf1, rentunf2	Dummy variables - housing tenure (have mortgage or partown, own outright, rent free, rent furnished, rent from Local Authority or Housing Association, other rent unfurnished)
totchild	Number of children in HH
totfemal	Number of adult females in HH
totmal	Number of adult males in HH
totother	Number of adults with employment status of sick/unemployed/unoccupied in HH
totretir	Number of retired adults in HH
totsick	Number of adults with employment status of sick in HH
totunemp	Number of adults with employment status of unemployed in HH

Cluster test labels

The cluster tests have coded names that relate to the variables and the method used to define them. For example, 42b.1 uses variable list 42b and cluster method 1.

Variables used in the test:

- 42b = number of adults, adult employment status, disaggregated housing tenure and number of cars used
- 41 = number of adult females, number of adult males, number of children, adult employment status, aggregated housing tenure, number of cars used and expenditure region category.

Cluster method code:

- 1 = systematic coding of household characteristics, sensible definitions that are easy to recreate
- 2 = optimal ordering produced by GCA is followed as closely as possible resulting in a number of clusters with long definitions that are not easy to recreate.

Appendix 3: Household income decile shares of expenditure by category

(a) FES

Expenditure category		Equivalent household income decile groups										Total
		1	2	3	4	5	6	7	8	9	10	
1	Household services etc	4.0	4.0	4.3	7.5	8.1	10.1	11.5	13.1	16.3	21.0	100.0
2	Motoring expenditure	3.1	2.2	2.8	4.9	6.1	9.3	15.0	13.3	18.7	24.8	100.0
3	Food (which attracts VAT)	4.3	4.7	4.9	7.1	8.3	10.4	12.1	13.3	16.1	18.8	100.0
4	Leisure goods and services	4.5	5.0	4.9	6.4	10.0	10.9	11.1	12.4	15.3	19.4	100.0
5	Adult clothing and footwear	3.9	4.1	4.7	6.3	7.6	9.7	12.7	14.2	15.4	21.4	100.0
6	Household goods	4.5	5.0	5.4	7.7	8.5	9.8	11.9	13.3	14.8	19.1	100.0
7	VAT-exempt goods	3.3	2.1	4.0	6.0	5.4	7.6	12.3	15.8	17.0	26.6	100.0
8	Food (zero-rated)	6.8	7.5	8.3	9.6	9.8	11.0	11.5	11.6	11.8	12.1	100.0
9	Books + newspapers	6.4	6.1	7.4	8.1	8.2	10.2	12.1	12.8	13.0	15.7	100.0
10	Domestic fuel and power	8.1	8.4	8.8	9.3	9.6	10.5	10.7	11.3	11.3	12.0	100.0
11	Other zero-rated goods	4.8	5.1	4.6	6.5	6.5	8.0	8.5	10.1	12.3	33.7	100.0
12	Children' s clothes	5.8	7.5	7.1	10.9	10.7	12.5	12.0	10.8	11.2	11.5	100.0
13	Insurance premia	4.1	3.5	4.5	6.2	7.6	9.8	12.7	14.3	16.4	20.9	100.0
14	Beer	3.9	5.2	5.7	6.6	8.2	10.8	11.9	14.7	17.1	15.8	100.0
15	Cider	5.1	8.5	4.8	9.7	10.5	9.4	12.4	14.4	14.2	11.0	100.0
16	Fortified wine	3.9	3.1	5.5	6.2	8.4	11.5	11.4	15.6	16.4	18.0	100.0
17	Wine	2.2	2.3	3.8	5.7	5.5	7.8	9.8	11.6	20.3	31.1	100.0
18	Champagne	4.0	0.8	3.8	1.6	6.0	7.6	11.6	10.7	20.4	33.5	100.0
19	Spirits	4.2	4.6	5.9	8.4	8.2	11.9	11.5	12.3	16.5	16.4	100.0
20	Cigarettes	7.3	10.6	10.2	10.4	11.0	11.2	11.6	10.4	10.2	7.1	100.0
21	Cigars	1.5	4.9	5.3	6.7	10.0	9.1	14.0	7.7	22.8	17.9	100.0
22	Pipe tobacco	10.3	11.9	10.7	11.1	14.3	12.1	10.8	9.0	4.8	5.0	100.0
23	Motor fuel	4.5	3.9	5.3	7.4	8.5	10.5	13.3	14.4	15.8	16.6	100.0
24	Motor fuel (diesel)	3.5	3.4	2.4	4.5	7.0	14.1	16.4	14.9	19.0	14.7	100.0
25	Pools stakes	5.6	8.3	7.9	10.8	10.1	12.8	12.9	12.6	10.4	8.6	100.0
26	Other betting stakes	5.4	9.8	7.4	11.9	11.4	14.1	8.4	9.7	10.4	11.4	100.0
27	Lottery stakes	5.8	6.6	7.8	9.0	10.1	11.9	13.6	12.6	12.7	9.9	100.0
	Total	4.7	4.9	5.5	7.4	8.1	9.9	11.9	13.2	15.0	19.4	100.0

(b) ratios FRS : FES

Expenditure category		Equivalent household income decile groups									
		1	2	3	4	5	6	7	8	9	10
1	Household services etc	1.22	1.19	1.19	0.87	0.93	1.01	0.99	1.03	0.96	0.97
2	Motoring expenditure	1.27	1.62	1.32	0.87	1.05	1.00	0.76	1.04	0.88	1.11
3	Food (which attracts VAT)	1.15	1.11	1.06	0.89	0.94	0.93	0.93	1.01	0.97	1.09
4	Leisure goods and services	1.11	1.00	0.97	0.98	0.82	0.87	1.12	1.09	1.02	1.01
5	Adult clothing and footwear	1.05	1.11	0.96	0.90	1.01	0.98	0.87	0.96	1.08	1.07
6	Household goods	1.17	1.15	1.04	0.88	1.02	1.01	0.95	0.95	0.98	1.03
7	VAT-exempt goods	1.27	1.69	0.95	0.85	1.17	1.18	0.91	0.88	1.11	0.91
8	Food (zero-rated)	1.01	1.02	0.92	0.90	0.99	0.95	0.96	1.02	1.04	1.14
9	Books + newspapers	0.96	1.10	0.92	0.94	1.06	1.02	0.93	0.97	1.02	1.06
10	Domestic fuel and power	1.03	1.07	0.95	0.96	1.02	0.97	0.98	0.98	0.99	1.05
11	Other zero-rated goods	0.98	0.91	1.08	0.83	1.03	0.93	1.08	1.01	1.00	1.03
12	Children' s clothes	1.04	0.94	0.93	0.69	0.93	0.92	1.07	1.21	1.12	1.12
13	Insurance premia	1.17	1.25	0.96	0.96	1.02	1.00	0.90	0.98	1.01	1.01
14	Beer	1.19	0.99	0.85	0.97	0.96	0.91	1.03	0.95	0.93	1.20
15	Cider	1.36	0.66	1.03	0.74	0.75	1.21	1.02	0.96	1.09	1.28
16	Fortified wine	1.32	1.63	1.04	0.91	0.96	0.82	1.02	0.93	1.04	0.99
17	Wine	1.85	1.38	1.01	0.88	1.17	1.02	1.02	1.18	0.88	0.90
18	Champagne	1.06	4.02	0.70	3.52	0.95	0.60	0.96	1.21	0.76	1.03
19	Spirits	1.04	1.31	1.05	0.90	1.15	0.81	0.98	1.08	0.89	1.06
20	Cigarettes	1.20	0.94	0.91	0.98	0.84	0.98	0.92	1.08	0.94	1.44
21	Cigars	2.03	0.68	0.54	1.02	0.96	1.08	0.68	1.81	0.89	1.15
22	Pipe tobacco	0.97	1.16	0.97	0.96	0.79	0.86	0.88	0.93	1.70	1.48
23	Motor fuel	1.01	1.23	0.93	0.87	0.97	1.01	0.92	0.98	1.05	1.07
24	Motor fuel (diesel)	1.43	1.05	1.81	1.22	1.10	0.58	0.83	0.92	0.99	1.33
25	Pools stakes	1.10	0.89	0.80	0.91	1.05	0.87	0.92	1.04	1.26	1.24
26	Other betting stakes	1.31	0.80	1.07	0.81	0.96	0.75	1.28	1.26	0.95	1.15
27	Lottery stakes	1.00	1.09	0.92	0.92	0.95	0.95	0.84	1.02	1.06	1.31
	Total	1.12	1.13	0.99	0.89	0.99	0.99	0.94	0.98	1.02	1.03

Notes: shading shows differences of 15% or more