WP. 27
ENGLISH ONLY

UNITED NATIONS STATISTICAL COMMISSION and ECONOMIC COMMISSION FOR EUROPE CONFERENCE OF EUROPEAN STATISTICIANS

EUROPEAN COMMISSION STATISTICAL OFFICE OF THE EUROPEAN COMMUNITIES (EUROSTAT)

Joint UNECE/Eurostat work session on statistical data confidentiality
(Bilbao, Spain, 2-4 December 2009)

Topic (v): Statistical disclosure control methods for the next census round

BALANCING RISK AND UTILITY – STATISTICAL DISCLOSURE CONTROL FOR THE 2011 UK CENSUS

Invited Paper

Prepared by Keith Spicer and Caroline Tudor, Office for National Statistics, United Kingdom

Balancing Risk and Utility – Statistical Disclosure Control for the 2011 UK Census

Keith Spicer[*] and Caroline Tudor[*]

[*]The Office for National Statistics, Segensworth Road, Titchfield, Fareham, PO15 5RR, UK, keith.spicer@ons.gov.uk caroline.tudor@ons.gov.uk

Abstract: We will describe the approach that has been adopted by the Office for National Statistics (ONS) to develop a statistical disclosure control strategy for 2011 UK Census tabular outputs. Focus will be given to the comparison of different disclosure control methods in terms of the protection they provide (risk) balanced against the impact that they have on the usefulness of the statistics (utility). We will describe the range of both qualitative and quantitative measures that have been used to evaluate both risk and utility and how these have been used to test methods on 2001 Census data. We will also discuss issues related to the practical implementation of an SDC method for Census tables such as the interaction with imputation for non-response, dissemination systems and table design.

1 Introduction

UK censuses are traditionally held every ten years and census outputs are the facts and figures produced from the questionnaires collected on census day. The information provided allows central and local government, health authorities and many other organisations to plan housing, education, transport and health services. Information from the questionnaires is processed to produce a database from which census outputs are created. Census output can be released in a number of different formats; standard pre-planned tables, commissioned tables requested by users, user defined tables via flexible table generating software and census sample microdata. Publishing aggregate or individual data carries the risk that individuals or entities could be identified and confidential information about them released. The UK Census Offices need to protect the confidentiality of census respondents for a number of reasons. The production and use of official statistics depends on the cooperation and trust of citizens. Such trust cannot be maintained unless the privacy of individuals' information is protected. There are also legal and policy obligations that must be respected. This paper will describe ongoing work on Statistical Disclosure Control (SDC) in preparation for the 2011 Census in providing a commitment to protecting confidentiality of census outputs.

In November 2006 the UK SDC Policy position (ONS (2006)) for the 2011 Census was agreed by the Registrars General of Scotland, England and Wales and Northern

Ireland. The Registrars General agreed to aim for a common UK SDC methodology for 2011 Census outputs to achieve harmonisation. Since agreeing the UK SDC Policy, the Statistics and Registration Service Act (SRRSA) 2007 came into force. The UK SDC Policy is in line with Section 39 of the SRSA. Following the introduction of the Act, the National Statistics Code of Practice has been superseded by the Code of Practice for Official Statistics (January 2009), which provides a similar level of confidentiality protection as the old code and hence the UK SDC policy is still valid.

An important step towards achieving harmonisation across the UK Census Offices was agreement of a SDC policy position setting out what constitutes a disclosive risk. The Registrars General considered that, as long as there has been systematic perturbation of the data, the guarantee in the Code of Practice would be met. It was therefore agreed that small counts (0's, 1's, and 2's) could be included in publicly disseminated Census tables provided that a) uncertainty as to whether the small cell is a true value has been systematically created; and b) creating that uncertainty does not significantly damage the data. The Registrars General expressed a preference for pre-tabular methods, provided there is not undue damage to the data, and have also stated that the key risk is attribute disclosure. Determining the exact threshold of uncertainty is an important part of the work and the Registrars General have reported the preference for SDC to have a ‘light touch’.

The aim of the SDC project is to design a UK SDC strategy in accordance with the agreed policy which protects against disclosure whilst maintaining maximum utility of the data. The strategy will cover (pre-defined) tabular outputs, microdata samples and possibly flexible user defined tabular outputs whilst taking into account the impact of interactions between these types of output. The strategy will also be designed to address the concerns of users whilst adequately protecting the data.

After thorough review and evaluation, a pre-tabular method of record swapping for disclosure control of census tabular outputs has now been signed off by the UKCC. The rest of this paper will describe the approach that has been adopted by the Office for National Statistics (ONS) in coming to a decision on this method. Focus will be given to the comparison of different disclosure control methods in terms of the protection they provide balanced against the impact that they have on the usefulness of the statistics. We will also discuss issues related to the practical implementation of an SDC method for Census tables such as the interaction with imputation for non-response, dissemination systems and table design. In addition some ideas for improving the methodology for record swapping in 2011 will be discussed.

2 Short-listed Methods

In July 2007 a review of a wide range of SDC methods was undertaken, assessing them against a set of qualitative criteria in line with the policy statement made by the Registrars General. Initially thirteen SDC methods were considered for short-listing and those did not meet certain criteria, such as not providing additive and consistent tables, were rejected outright. Longhurst et al. (2007) describe this work in detail. The process resulted in three SDC methods being short-listed for further evaluation:

- Record swapping
- Over-imputation
- ABS Cell Perturbation Method (developed by Australian Bureau of Statistics)

2.1 Record Swapping

Record swapping is a well-known SDC method which involves swapping the geographical identifiers of a small percentage of household records with other records matching on specific control variables. Matching on control variables helps to minimise bias. Record swapping in the UK would generally be carried out within a Local Authority District (LAD)[1] and households / persons are swapped in and out of smaller geographical areas e.g. Output Areas. If a match can be found for every record selected, record swapping therefore ensures that Local Authority marginal distributions remain unaffected[2]. Record swapping can take different forms including random record swapping or targeted record swapping. Random record swapping involves selecting, at random, households and individuals for swapping. Targeted record swapping involves selecting a random sample of the potentially unique/risky records for swapping. In our evaluation, a random sample within strata defined by control variables was selected using a fixed swapping rate. Control variables of hard-to-count index[3], household size, sex and broad age distribution of the household were used. All geographical variables were swapped. For the targeted swap, high risk records (defined as those contributing to small cells in a set of tables) were identified and flagged. The swap was then implemented by pairing and swapping households that matched not only on the control variables but also on the flagged variable.

---

[1] In 2001, Output Areas were the smallest geographic building brick, that could be combined to form higher geographies such as Local Authorities.

[2] In practice, there is a small amount of swapping between LADs since it is otherwise difficult to find matches for very unusual households

[3] The Hard to Count (HtC) index was constructed in the 2001 UK Census as a measure of enumeration difficulty. It was constructed from the following 1991 Census variables; Multi-occupancy, unemployment, language difficulty, private rented accommodation, number of household imputed in 1991. Scotland also used ethnic group.

2.2 Over-imputation

Imputation is a commonly used method for replacing missing values in census and survey data due to item non-response. A new method of over-imputation was devised for disclosure control of census data using CANCEIS (developed by Statistics Canada – Canadian Census Edit and Imputation Software). In the earlier stages of the evaluation, work concentrated on geographic imputation, to attempt for consistency with record swapping. Geography and age were blanked out and imputation used as a method of disclosure control, i.e. over-imputation. Donors from the remaining population were used to replace values, in the case of imputing geography; OAs were imputed within the same LAD, whereas age was imputed using all possible donors. However it was felt that this method of imputation didn't exploit the full benefits it could offer over record swapping so a later evaluation looked at imputing variables other than geographic ones. Non-geographic over-imputation was thought likely to provide better protection for both tables and microdata. Donors are found from the remaining population that match or closely match on other related variables. CANCEIS will attempt to find the best possible match, and then failing that on a subset of them. CANCEIS might not necessarily match on all variables for a given recipient/donor pair, but uses a probabilistic method of choosing which combination of matching variables is best in a given case. There is also a distance function specified to find or prioritise donors within certain geographical limits of the recipient. As with record swapping, over-imputation was applied to the random sample of records and then to the specific high risk records for targeted imputation.

2.3 ABS Cell Perturbation

This new cell perturbation method developed by the Australian Bureau of Statistics (ABS) is essentially a post tabular approach which takes into account pre-tabular information (Fraser and Wooton, 2006). The method involves adding small perturbations[4] to all cells in a table using a two stage process. Stage one results in a consistently perturbed non additive table. To achieve this, all microdata records are assigned a record key. When creating a table the record keys for all records contributing to each internal cell are summed and a function is applied to this sum to produce the cell key. Lookup tables (determined by the organisation) are then used where the true cell value and the cell key are used to determine the amount by which the cell count should be perturbed. This means that the same cell is always perturbed in the same way. The perturbation can be set to zero for a pre-determined set of key outputs (e.g. age by sex population counts). Table margins are perturbed independently using the same method. At the second stage another perturbation is added to each cell (excluding the grand total) to restore table additivity. The stage two perturbations are generated using an iterative fitting algorithm which attempts to balance and minimise absolute distances to the stage one table, although not

[4] Note that some perturbations will be zero

necessarily producing an 'optimal' solution. For this evaluation a modification of the original ABS method, referred to as Invariant ABS Cell Perturbation (IACP), has been developed to improve utility by attempting to make the first stage perturbations invariant with respect to the table cell frequencies[5]. Details of this method are described in Shlomo and Young (2008). Implementation of the IACP method involved assigning microdata keys to the individual records and then perturbation applied once the table had been created.

2.4 Small Cell Adjustment (SCA)

Applying small cell adjustments involves randomly adjusting small cells upwards or downwards to a base using an unbiased prescribed probability scheme. Marginal totals are obtained by summing perturbed and non-perturbed cells. Small cell adjustments were used in addition to random record swapping to protect 2001 Census tabular outputs for England and Wales and Northern Ireland. The method of SCA plus record swapping was used as a base for comparison in our evaluation.

3. Risk-Utility Evaluation of Short-listed methods

The most important characteristic of the SDC strategy for 2011 Census is that disclosure risk should be managed to an acceptable level, in order to respect legal and policy obligations and to ensure public co-operation and trust are maintained. Census outputs have a higher risk of disclosure and are harder to protect than other statistical data outputs because they contain whole population counts, because small areas predominate in output geography, and because tables are disseminated from only one data source, therefore meaning tables can be linked and differenced.

The Registrars General have highlighted that the key disclosure risk for 2011 Census output is attribute disclosure, which is highly associated with – but not exclusively so – low numbers in tables. Hence the analyses in this paper concentrate on the effects of possible SDC methods on the numbers of zeroes, ones and twos in tables. Managing risk will necessarily impact on data utility. The trade off between risk and utility is evaluated quantitatively. Unperturbed 2001 Census microdata were obtained for two Estimation Areas (EA) in England & Wales; SJ (Southampton, Eastleigh and Test Valley districts) and KB (Congleton, Chester, Crewe and Nantwich, Ellesmore Port and Vale Royal districts). KB (523,465 persons, 215,869 households) is a rural area chosen for the sparsity of its population whereas SJ (437,744 people, 182,337 households) is more urban and densely populated. Only a brief summary of the results are presented in this paper as a very extensive evaluation was carried out by the SDC team. The following set of tables was analysed:

[5] There may be a small increase in the number of zeros in tables when applying IACP

(CTable 1) Country of Birth (2) by Sex (2) by Religion (8) by ward
(CTable 2) Number of persons in household (4) by Accommodation Type (3) by OA / ED
(CTable 3) Age (16) by Sex (2) by Marital Status (2) by OA / ED
(CTable 4) Origin-Destination Flows from OA / ED to TTWOA
where TTWOA is travel to work OA for all in England and Wales.

Three further census tables were analysed with the purpose of assessing more specific features of the SDC methods particularly additivity, consistency and disclosure by differencing:

(CTable 5A) Age (9) by Ethnic group (17) by Sex (2) for all persons in wards
(CTable 5B) Age (9) by Ethnic group (17) by Sex (2) for persons without limiting long term illness in wards
(CTable 5C) Age (9) by Ethnic group (17) by Sex (2) for all persons with LLTI in wards

CTable 5B relates to a subpopulation of CTable 5A. CTable 5C is obtained by differencing CTable 5B from CTable 5A. For our evaluation CTable 5C was produced as a disclosure-controlled table independent from CTables 5A and 5B and as a derived table via the two disclosure-controlled CTables 5A and 5B. The latter scenario represents a situation common to the 2001 Census where census users made 'special requests' for specific tables which were very similar in composition and differenced tables could be produced indirectly. The tables vary in terms of number of small cells, zeros and average cell size. Different levels of perturbation were studied, ranging from 2% to 20%.

3.1 Disclosure Risk Measures

Disclosure risk was evaluated by assessing group disclosure, within-group disclosure, negative attribute disclosure and an analysis of the small cells in the table; looking at the proportion of small cells left unchanged after protection had been applied. Group disclosure was a measure of all respondents fall into a single response category for a particular variable for a given category of the other variable. Within-group disclosure looked at cases where responses were spread across two categories and one of these categories only contains one person. Negative attribute disclosure examined cases where no responses fell into a row or column. Fig 1 shows one example from the evaluation examining group disclosure indicating how, at the 2% level, record swapping when targeted to risky records performs best, although protection is also afforded by imputation and IACP. Note that this is a fairly low level of perturbation so most instances of group disclosure remain.

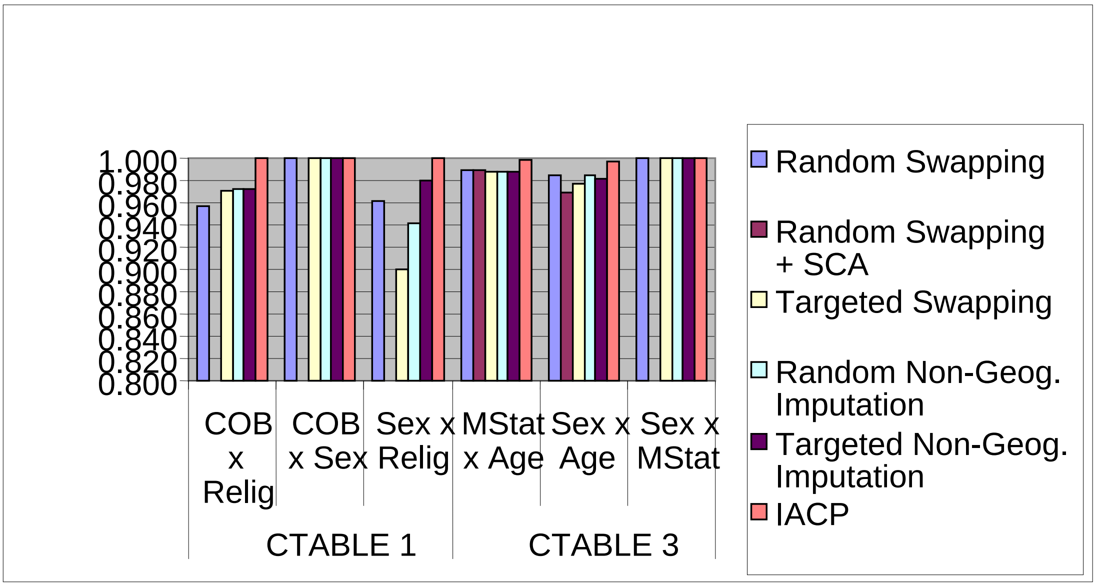

Fig 1. Proportion of columns with group disclosure remaining in protected table.

Data on within-group disclosure and negative attribute disclosure showed a similar pattern amongst the SDC methods. The level of perturbation used in the evaluation for all of the methods provides a small amount of protection but consideration would need to be given as to what level provides sufficient uncertainty. In practice, some protection is provided through user perception, given that there has been some SDC employed, without users being aware of the full details. In 2001, far more uncertainty was added to the data through the edit and imputation procedures used to counter item and person non-response and inconsistent answers to census questions than through disclosure control and this is likely to be the case too in 2011.

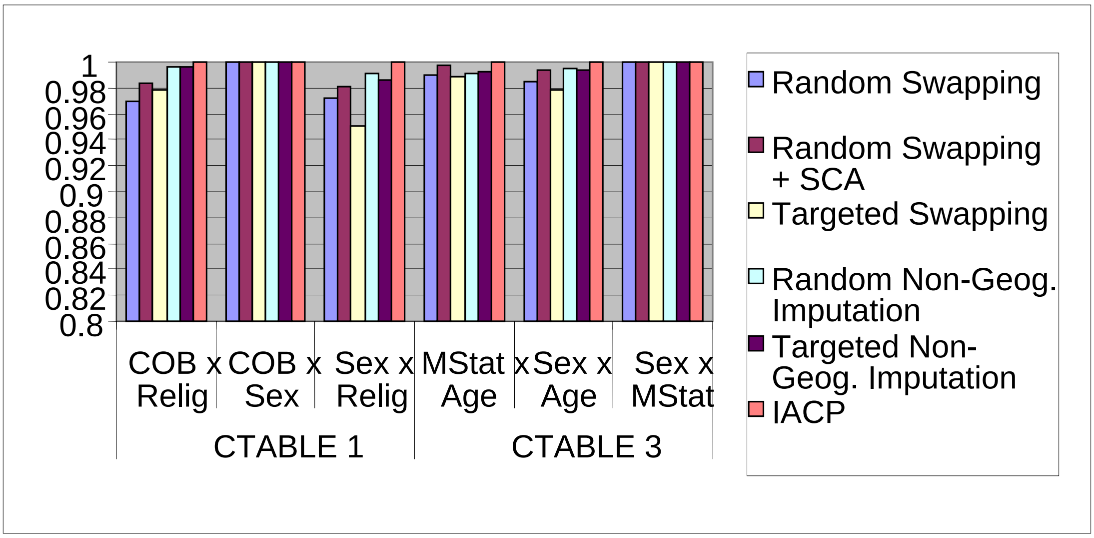

Fig 2. Proportion of zeros in the raw table left unchanged under each evaluated method

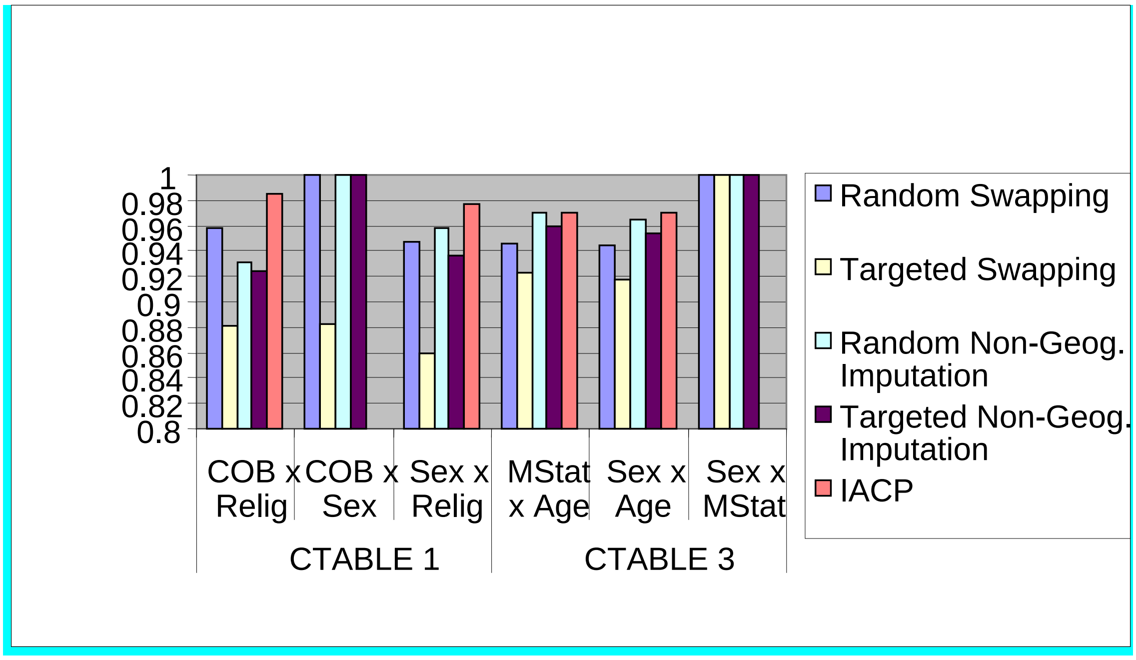

Fig 3. Proportion of cells of size 1 or 2 remaining in the table, across protected methods.

Fig 2 and Fig 3 show the disclosure risk remaining after the SDC methods have been applied. IACP appears to leave the highest number of small cells not being perturbed, i.e. leaving the most risk but it is not directly comparable to the other methods since the perturbation levels relate to cells rather than records. For both CTable 1 and CTable 3, imputation does perturb more small cells than IACP but swapping appears to protect more. Targeted swapping does perturb the highest percentage of small cells. The effect of using targeted rather than random methods is to protect more cases of small cells, as one would expect, though the extent of the difference between random and targeted methods is a bit mixed.

All methods provide some protection for disclosure by differencing. Pre-tabular methods will protect in the same way that they protect against other forms of disclosure, by perturbing the microdata so that when producing two tables and subtracting, the resultant table will be identical to that produced by interrogating the microdata directly for the ‘sliver’. Hence the protection will be equivalent to that given to the records in the sliver. On considering the post-tabular method, IACP, CTable 5C has been produced in two ways: (i) by differencing between two tables (5A and 5B) protected independently and (ii) by constructing directly from the microdata and then protecting., This does give rise to a small number of cells where differencing causes apparent negative cell counts to appear. For instance, if CTable 5A and 5B have true cell counts of 6 and 5 (so the true cell count for the difference is 1), they could be perturbed such that the cell count for 5A is less than that for 5B, say 4 and 5.

4.2 Utility measures

The in-house Information Loss Software[6] was used to evaluate the information loss associated with the short-listed SDC methods and the measures used are described in Shlomo and Young (2006). The software calculates a variety of information loss metrics by comparing the protected data with the original pre-disclosure controlled data. These include measurement of change in variance across rows of the table, change in level of association in the table, distance metrics, impact on totals and subtotals, change in rankings within a variable. As well as these properties we also looked at consistency and additivity in the protected tables. As before we present a short selection of the analyses for illustrative purposes.

It should be noted that for many of the variables the level of association is low. We can also take the counts at ward level (in the case of CTable 1) and OA level (CTable 3), and assess the changes in association at the lower geographies. Those results are shown in Table 1. Record swapping generally performs better than over-imputation though there is some effect on associations at the lower geographies, since records

[6] Infoloss software package has been developed in-house in SDC Methodology.

are swapped between wards and OAs, the geographies at which the associations are being assessed. IACP has the least effect on associations while the 2001 method of record swapping plus small cell adjustment has greatest effect, this being most marked where the tables are sparser.

	CTABLE 1			CTABLE 3		
% change	COB x Relig	COB x Sex	Sex x Relig	MStat x Age	Sex x Age	Sex x MStat
Random Swapping	0.15	-0.93	0.17	-0.03	-0.15	-0.12
Random Swapping + SCA	1.04	-0.93	2.01	1.29	6.95	-0.12
Targeted Swapping	0.04	-2.74	-0.35	0.00	0.19	0.06
Random Non-Geog. Imputation	0.89	0.37	5.26	-0.01	0.01	0.85
Targeted Non-Geog. Imputation	0.26	0.59	5.15	0.04	0.26	0.43
IACP	-0.05		0.16	-0.00	0.00	0.02

Table 1. Changes in Association (Cramer's V) for evaluated methods and tables.

Table 2 examines Relative Absolute Deviation. The Relative Absolute Deviation (RAD) values are expressed in percentage terms, so that, for example, on average, across the COB * religion cells in CTable 1 the values are changing by 0.36 per cent.

	CTABLE 1			CTABLE 3		
	COB x Relig	COB x Sex	Sex x Relig	MStat x Age	Sex x Age	Sex x MStat
Random Swapping	0.36	0.02	0.35	0.71	0.68	0.02
Random Swapping + SCA	2.27	0.32	2.23	5.16	4.39	0.05
Targeted Swapping	0.71	0.05	0.60	0.85	0.84	0.02
Random Non-Geog. Imputation	0.85	0.02	0.85	0.41	0.49	0.02
Targeted Non-Geog. Imputation	0.90	0.02	0.90	0.49	0.60	0.02
IACP	0.06		0.12	0.38	0.38	

Table 2. Relative Absolute Deviation (RAD) for evaluated methods and tables.

The largest percentage changes are incurred by using the SCA method, where small cells will be adjusted, the relative change perhaps a doubling or trebling of the cell count. IACP performs best here, having the lowest relative change, better than both the pre-tabular methods, whether random or targeted. Generally, imputation performs slightly better than swapping except in those tables with the religion variable.

Other desirable properties of an SDC method are consistency and additivity. Since record swapping and over-imputation are pre-tabular methods and involve perturbation of the microdata before tables are created these methods will always produce consistent tables, i.e. the same cell in a different table will always have the same value. The use of microdata keys in the IACP method ensures that the same cell has the same pertubation each time it falls in a table, however some of this consistency is lost when the table is made additive. Differencing tables protected using the IACP method may also produce inconsistencies. This was tested using CTables 5A-C. CTable 5C was produced by differencing CTable 5A and 5B after they had been protected using the IACP method.

	Difference between CTable 5C and 5C*			
	Total no. differing cells	+/-1	+/-2	>2
IACP 98%	468 (1.2%)	426 (1.1%)	31 (0.08%)	11 (0.03%)
IACP 90%	1,555 (3.8%)	758 (1.8%)	725 (1.8%)	72 (0.2%)
IACP 80%	2,443 (5.9%)	1,871 (4.5%)	383 (0.9%)	189 (0.5%)

Table 3. Number (and percentage) of Cell counts differing between CTables 5C and CTable 5C*

CTable 5C is derived by the difference between the protected CTable 5A and the protected CTable 5B

CTable 5C is derived by protecting the table produced by subtracting the unprotected CTables 5A and 5B*

The results in Table 3 show that the IACP method does not maintain consistency between differenced tables. For example for IACP 98% there are 468 cells whose values differ between the two tables, the majority of these differ by an absolute value of 1. The maximum difference is 5.

In terms of additivity, since record swapping and over-imputation are pre-tabular methods and involve perturbation of the microdata rather than the tables themselves, these two methods necessarily preserve additivity in all tables. The IACP method is post-tabular and involves perturbation of table cells, so additivity is not preserved in all cases. However, the IACP algorithm restores the additivity. Hence all three short-listed methods would produce outputs that are additive.

4.3 Balancing Disclosure risk and utility

Graphs were plotted of disclosure risk measures against a variety of utility measures. Ideally, we would like to compare the utility across the different methods where the risk is approximately the same. As an example, Fig 4 uses the proportion of group disclosure remaining in the protected tables as a measure of disclosure risk, and is compared to the ratio of variances (across rows of the table, measured before and after SDC).

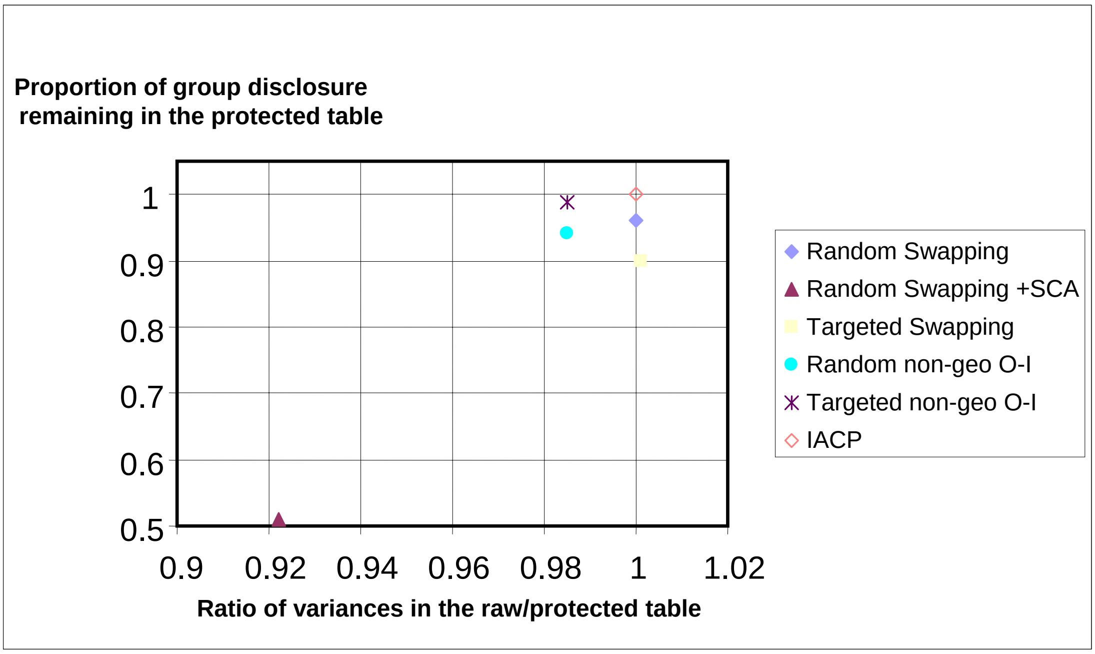

Fig 4. Proportion of group disclosure remaining in the protected table and the ratio of the variance change between the raw and protected tables (using CTable 1)

The 2001 method of random swapping with SCA has the largest effect on the variance (reducing variation in protected tables) – see Fig 4. But this method also reduced the group disclosure in the protected tables the most. The other evaluated methods slightly reduced the variances, and this was more noticeable with over-imputation methods. Random and targeted swapping had similar levels of group disclosure remaining in the tables, performing slightly better than over-imputation. A full evaluation report with all the risk and utility measures, in combination with risk-utility plots can be obtained on request team at ONS.

4.4 Origin-Destination Tables

Origin destination (O/D) tables are different to census area statistics tables in that they consist of data for all combinations of areas in England and Wales (in each O/D table), and can have over 10 million cells depending on the breakdown of variables and level of geography. Zeros typically comprise 98-99% of the table cells at OA level with small cell values making up the majority of non-zero cells though zero rows are usually suppressed from output. Determining an appropriate SDC strategy for O/D tables is very problematic due to their sparsity. Post-tabular methods showed to provide some protection but have a significant impact on data utility (at low geographical levels) since flows will disappear from the table. For the pre-tabular methods it is likely that illogical flows will occur in the protected table, i.e. walking 60 miles to work. The recommendation was made that protection for O/D tables (particularly at the low geographical levels) should be provided by licensing and restricted access. At higher levels an SDC method could be applied or it may be determined that no additional protection (other than aggregation) is required since the flows are not so disclosive (this will depend on variable breakdowns).

5 Additional Evaluation

The results from the risk-utility evaluation were not always clear cut and no one SDC method outperformed all others. It is important to note that this work was quality assured by academics who specialise in SDC and census experts from other NSIs, in the form of the UK CDMAC SDC subgroup which reviewed the work at several stages giving suggestions for improvement. To assist in making the final decision a set of scoring criteria were defined. Several key criteria were labelled as mandatory. There was a requirement for the method to satisfy or mostly satisfy all mandatory criteria in order to be considered for the final method. The criteria took into consideration:-

MANDATORY CRITERIA

- The method creates the desired level of doubt about any attribute disclosure and protects against differencing
- Marginal totals in protected tables are unbiased
- Protected tables are additive
- The method cannot be unpicked

SECONDARY CRITERIA

- Method provides consistent cell counts and totals between different protected tables
- The method is practical bearing in mind the resources available in terms of manpower, computing power and software costs.
- For a given level of risk relationships between variables are maintained in protected tables.
- The method can take into account the levels of imputation and overall data quality of different variables
- Counts of households and residents for small areas are not unduly perturbed
- The method does not unduly perturb/affect counts for large geographies (e.g. LA level and above)
- The method has a low impact on the variance of estimates
- The method can be used or adapted to protect outputs from special populations such as communal establishments or from workplaces
- Will not restrict the detail of releases or the subsequent protection method to be used for microdata samples
- The method and any required software will have adequate lifespan for purpose

- The method can easily be accounted for by users in analysis
- The same method can be applied to microdata outputs
- The method is likely to be easily understood by users
- The method has been effectively used for protecting similar outputs
- The method makes use of all data collected in the Census
- The method will be applied systematically to all tables and all cells

The IACP method was rejected. The main disadvantages in comparison to the pre-tabular methods were that: the method did not maintain consistency between tables, the feasibility and practicality of implementing the method are unknown and questioned, the method is not easy for users to understand and the method cannot take into account imputation levels already present in the data. Moreover the Registrars General had expressed a preference for a pre-tabular method.

The scoring results showed there was little to choose between the other two short-listed methods. The main differences between the two methods and their relative advantages and disadvantages (based on the criteria) were: both methods are able to take into account levels of imputation already present in the data but over-imputation can do this at a variable level, record swapping does not distort relationships between variables (because only geography is swapped), the results show that over-imputation does. In fact this method impacts on totals and subtotals across tables, e.g. after protection the table may contain more men and less women than the original table, above local authority level record swapping has a minimal impact on the utility of the tables. This however, could cause problems if users matched higher level tables against microdata samples and determine and locate population uniques. However, this risk could be managed by licensing arrangements. Record swapping was the preferred method and this has been agreed and signed off by UK Census Committee.

5. Practical Implementation

Further work is now continuing to determine the exact methodology for record swapping. A chief aspect of this work is assessing how to target risky records. There has been a precedent for a targeted approach in the US Censuses 2000 and planned for 2010 (Zayatz, 2009). Similar approaches may be considered for the UK Census where applicable. Since the main risk is attribute disclosure, a “key” formed of a set of tables potentially leading to attribute disclosure is suggested to identify risky records. The idea would be that those households unique across the greatest number of tables would be best targeted for swapping. Preliminary investigations have suggested that households unique on one table tend to have characteristics that lead to them being unique on other tables. In addition, certain households/records would be selected as special cases and given a high probability for swapping, for example households with an ethnic make-up unique in their OA.

An advantage of pre-tabular methods is that one can take into account imputation for non-response when implementing the method. An argument often raised by users is that disclosure control is unnecessary in some areas where response to the census is poor. Where there is low response, persons and households will have been imputed through the equivalent of the 2001 One Number Census. There will also be item imputation where an item is missing or is found, subsequent to capture, to be inconsistent with other variables in the individual or household records. It will be possible to take into account levels of imputation when carrying out targeted swapping by limiting the amount of perturbation in these areas.

Moreover the selection of risky records, from those identified using the key, would be stratified by OA to give a more equal chance of introducing uncertainty about all cells. Other ideas include a swapping rate in inverse proportion to OA size (similar to the US approach). A new idea is also to distinguish between types of unique; those which are "local unique" and occurring at random due to small/sparse populations as compared to those which are special unique (unusual characteristics irrespective of geography) and therefore using a graduated approach to swapping in relation to distance swapped and degree of matching. Moreover it is proposed to have higher rates of swapping in slivers created via differencing although these slivers would need to be anticipated in advance.

This work will feed into setting rates for swapping; we will be considering the rates at different geographic levels, and swapping just enough records to create the 'sufficient uncertainty' desired by the RsG (and indeed to assess what that level of sufficient uncertainty is). The next stage will be to carry out a simulation study to test feasibility of the ideas proposed. Note that over and above the method selected, additional rules will be needed for thresholds and sparsity in tables. Work is currently taking place on differencing between geographies, to ascertain whether it is possible to output on both Census (super output areas) and administrative (ward) geographies without increasing the disclosure risk unreasonably.

A considerable amount of work will need to be undertaken to address how swapping with interact with output design and the additional rules that will be required around table design, particularly given plans for flexible table generation, in terms of restricting number of variables and determining minimum cell sizes.

In light of the choice of record swapping as the primary strategy for tabular outputs, we can now consider the protection needed for microdata products. Work on this is at a fairly early stage, but products are being considered at different levels of access, including at EUL and VML. This time, microdata will be an integral part of the outputs package, as opposed to after the main outputs, which was the case in 2001).

7 Summary

This paper has given an overview of the strategy leading to the decision of record swapping for protecting Census 2011 tabular outputs. The work that has informed this decision includes a thorough evaluation on Census 2001 data considering three short-listed methods in terms of risk and utility. Further work is continuing to determine the specifics of the swapping methodology and the interaction with outputs, table design, thresholds and imputation. An important point is that the record swapping methodology employed in 2011 will be different to that in 2001. It will be 'smarter', targeting risky records and using swapping levels appropriate to characteristics of the geographical area – size and level of imputation (non-response). So the additional protection that was afforded in 2001 by small cell adjustment is supplied in 2011 by undertaking record swapping in a 'smarter' way, retaining as much data utility as possible.

References

Longhurst, J., Tromans, N., Young, C. and Miller, C. (2007) Statistical Disclosure Control for the 2011 UK Census, Joint UNECE/Eurostat conference on Statistical Disclosure Control, Manchester, 17-19 December.

Fraser, B. and Wooton, J. (2006) A proposed method for confidentialising tabular output to protect against differencing, Internal report, Data Access and Confidentiality Methodology Unit, ABS.

ONS (2006) www.ons.gov.uk/census/2011-census/produce-deliver-data/regrs-gen-agreement.pdf

Shlomo, N. and Young, C. (2008) Invariant Post-tabular Protection of Census Frequency Counts, Proceedings of Privacy in Statistical Databases, Istanbul, Turkey, September 24-26.

Shlomo, N. and Young, C. (2006) Statistical Disclosure Control Methods through a Risk - Utility Framework: Proceedings of the Privacy in Statistical Databases CENEX-SDC Project International Conference, Rome, 13-15 Dec 2006.

Zayatz, Lucero, Massell, Ramanayake (2009 – forthcoming) Disclosure Avoidance for Census 2010 and American Community Survey Five-Year Tabular Data Products