

Working Paper M09/14

Methodology

Assessing The Disclosure

Protection Provided By

Misclassification For Survey

Microdata

Natalie Shlomo And Chris Skinner

Abstract

Government statistical agencies often apply statistical disclosure limitation techniques to survey microdata to protect confidentiality. There is a need for ways to assess the protection provided. This paper develops some simple methods for disclosure limitation techniques which perturb the values of categorical identifying variables. The methods are applied in numerical experiments based upon census data from the United Kingdom which are subject to two perturbation techniques: data swapping and the post randomisation method. Some simplifying approximations to the measure of risk are found to work well in capturing the impacts of these techniques. These approximations provide simple extensions of existing risk assessment methods based upon Poisson log-linear models. A numerical experiment is also undertaken to assess the impact of multivariate misclassification with an increasing number of identifying variables. The methods developed in this paper may also be used to obtain more realistic assessments of risk which take account of the kinds of measurement and other non-sampling errors commonly arising in surveys.

Assessing the Disclosure Protection Provided by Misclassification for Survey Microdata

Natalie Shlomo and Chris Skinner

University of Southampton

Abstract:

Government statistical agencies often apply statistical disclosure limitation techniques to survey microdata to protect the confidentiality of respondents. There is a need for valid and practical ways to assess the protection provided. This paper develops some simple methods for disclosure limitation techniques which perturb the values of categorical identifying variables. The methods are applied in numerical experiments based upon census data from the United Kingdom which are subject to two perturbation techniques: data swapping (random and targeted) and the post randomisation method. Some simplifying approximations to the measure of risk are found to work well in capturing the impacts of these techniques. These approximations provide simple extensions of existing risk assessment methods based upon Poisson log-linear models. A numerical experiment is also undertaken to assess the impact of multivariate misclassification with an increasing number of identifying variables. It is found that the misclassification dominates the usual monotone increasing relationship between this number and risk so that the risk eventually declines, implying less sensitivity of risk to choice of identifying variables. The methods developed in this paper may also be used to obtain more realistic assessments of risk which take account of the kinds of measurement and other non-sampling errors commonly arising in surveys.

Keywords and phrases: disclosure risk, identification risk, log linear model, measurement error, post randomization method, data swapping.

1. Introduction

Government statistical agencies have statutory and ethical obligations to protect the confidentiality of the data they collect. At the same time, their core mission is to ensure that these data are used effectively for statistical purposes. Tensions between these two objectives may arise, in particular, when access to microdata on individuals or establishments is to be provided to researchers, so that they may conduct their own analyses of social or economic phenomena. Although microdata may be anonymised by removing

obvious identifying information such as name and address without damage to the statistical analyses, such anonymisation will rarely be considered sufficient for confidentiality protection, since the rich socio-economic information in the microdata may often enable records to be identified by matching to another data source on known individuals or establishments. Agencies have therefore developed a number of ways of protecting confidentiality in this context. One common approach is to modify the microdata file by applying a statistical disclosure limitation (SDL) method, such as recoding or data perturbation, to those variables judged potentially identifying (Federal Committee on Statistical Methodology, 2005). Such modification can, however, seriously reduce the utility of the microdata and it is therefore important for the agency to be able to assess the protection provided by such methods in order to be able to make judgements about the degree of modification to apply.

The aim of this paper is to develop methodology to assess the disclosure protection provided by the misclassification of one or more categorical identifying variables. Misclassification is supposed here to arise in one of two ways. First, it may be the result of the deliberate application by the agency of an SDL method, specifically we consider the methods of data swapping (Dalenius and Reiss, 1982) and post-randomisation or PRAM (Gouweleeuw et al., 1998). This paper is motivated by experience of the use of such methods at government statistical agencies (especially in the United Kingdom) with microdata from social surveys on individuals or from population censuses. In these cases, the potential identifying variables which might be used for matching are invariably categorical. A second way in which misclassification may arise is as a result of measurement error which arises naturally in surveys and takes the form of misclassification for categorical variables (Kuha and Skinner, 1997). In this case, we shall suppose that the agency has some information about the nature of the misclassification mechanism.

In the current practice of statistical agencies, when the disclosure protection of such methods is assessed, it is usually based upon simple measures, such as functions of the diagonal elements of the misclassification matrix (Willenborg and De Waal, 2001, p119), or a simple estimated probability that an apparent match is correct (Gouweleeuw et al., 1998), or via the outcome of a record linkage experiment (see below). Reiter (2005) developed a more sophisticated approach by defining a measure of identification risk, based upon the modelling framework of Duncan and Lambert (1989), and showing how it could be assessed before and after the application of a number of SDL methods, including data swapping. This focus on identification risk is often appropriate in government contexts, where judgments about

protection are informed by legislation or codes of practice which express threats to confidentiality in terms of individual respondents being identified. However, the need to model a very wide range of microdata variables and relationships in Reiter's (2005) approach may limit its application in practice. In this paper we develop an approach which is based on a similar framework to Reiter (2005), but which retains some of the simplicity of the former methods. We achieve simplification by restricting the information set upon which the risk measure is conditioned, extending the approach of Skinner and Shlomo (2008). Our approach also extends Reiter (2005) by taking fuller account of the protection achieved from sampling.

Assessing identification risk using record linkage experiments (e.g. Yancy, Winkler and Creecy, 2002; Domingo-Ferrer and Torra, 2001) is natural given the threat that such methods pose (Fienberg, 2006). The experiment typically involves matching records in the microdata file, masked by an SDL method, to records in the original unmasked file. The risk is often defined as the proportion of such matches which are correct (Spruill, 1982). A problem with this approach is that it usually does not take account of the disclosure protection provided by sampling. We shall show in the Appendix that our proposed approach to assessing identification risk in the case of exact matching does, in fact, provide a closed form expression for the correct match proportion which would be estimated by an experiment using a form of probabilistic record linkage proposed by Fellegi and Sunter (1969). Record linkage experiments have the potential to capture the impact of a wider range of types of potential attack, including those that make explicit allowance for data masking and exploit greater computational power (Winkler, 2004), but consideration of such extensions is beyond the scope of this paper.

Statistical modeling approaches to identification risk assessment have been proposed by a number of authors (e.g. Paass, 1988; Duncan and Lambert, 1989; Fuller, 1993). It is generally assumed that a hypothetical intruder seeks to identify an individual in the microdata by matching records to known individuals in the population using identifying variables, also called key variables, values of which are known both for the microdata records and for the known individuals. This paper builds on the literature which has used models for categorical key variables as a basis for assessing disclosure risk. Bethlehem et al. (1990) is a seminal contribution. We follow especially Skinner and Shlomo (2008), who considered the use of log-linear models to assess identification risk. Their work did not, however, consider the impact of SDL methods on risk, other than the recoding of key variables.

The empirical work in this paper is based upon the 2001 population census

in Great Britain, which will be used to provide population data to validate risk assessments for samples, viewed as representing potential sample surveys. Our focus will be on the impact of SDL methods on identification risk. The effects of these methods on the utility of potential data analyses is also vitally important and we provide some information loss measures to analyse and compare the perturbation methods.

Our paper is organized as follows. Measures of identification risk in the presence of misclassification are developed in Section 2. Since these measures depend upon population quantities which may be unknown, methods of estimating these measures using sample data are considered in Section 3. Applications using census data are presented in Section 4 for a random and targeted data swapping method and a random and targeted post-randomization method (PRAM). A further numerical illustration with multivariate misclassification is presented in Section 5. Section 6 contains a concluding discussion.

2. Identification Risk Under Misclassification

Consider the release of a microdata file consisting of records for a sample $s = \{1, 2, \dots, n\}$ drawn from a finite population U of size N . We suppose an intruder seeks to match a known target unit in U to a record in the file using C categorical key variables X^1, \dots, X^C . We assume the agency knows the intruders choice of key variables. Possible departures from this assumption are discussed in section 6. The variable formed by cross-classifying the key variables, as measured by the intruder on the target unit, is denoted X and its values are labelled $1, 2, \dots, K$. The value of X recorded in the microdata, after the application of the SDL method (and natural measurement error) is denoted \tilde{X} . We treat the values of X for population units as fixed and suppose the values of \tilde{X} for the records in the microdata are determined independently by a misclassification matrix M , where

$$(2.1) \quad Pr(\tilde{X} = j | X = k) = M_{jk}.$$

To assess the disclosure protection provided by misclassification, we imagine that the intruder observes a match between a specific sample unit A and a target population unit B , i.e. observes $\tilde{X}_A = X_B$ (where \tilde{X}_A is the value of \tilde{X} for unit A and X_B is the value of X for unit B), and measure disclosure risk in terms of the uncertainty as to whether $A = B$. A simple ad hoc measure of this uncertainty is given by M_{jj} (or $1 - M_{jj}$), where j is the common value of \tilde{X}_A and X_B . Willenborg and De Waal (2001, p121) propose that the agency specifies upper bounds for these diagonal elements of M according

to the level of protection required. Following Reiter (2005) we define the identification risk as $Pr(A = B|data)$, where the values \tilde{X}_A and X_B are implicitly included in the *data* and the nature of the probability mechanism will be clarified later. A simplified approach to estimating this risk is given by Gouweleeuw et al. (1998), who make the very conservative assumption that the intruder knows that B is in the sample and approximate $Pr(A = B|data)$ by $Pr(X_A = j|\tilde{X}_A = j) = M_{jj}Pr(X_A = j)/\sum_k M_{jk}Pr(X_A = k)$, which they estimate by:

$$(2.2) \quad M_{jj}f_j / \sum_k M_{jk}f_k,$$

where f_k is the number of units in s for which $X = k$ (they in fact use the odds rather than the probability). In contrast to the highly simplifying assumptions of Gouweleeuw et al. (1998), Reiter (2005) allows for considerable generality by adopting a very wide definition of *data* in $Pr(A = B|data)$, so that it may include all the values of \tilde{X}_i in the sample as well as the values of any other microdata variables. This creates not only a major modelling task to assess the probability of interest, but also the possibility that this probability will be sensitive to the specification of the model. We seek an intermediate position, avoiding the very conservative assumption that the intruder knows that B is in the sample, but reducing the scope of *data* in $Pr(A = B|data)$ to avoid the complex modelling issues. We define the matching variable \tilde{Z}_i to be 1 if $\tilde{X}_i = X_B$ and 0 otherwise and we take the data to consist of the values \tilde{Z}_i for $i \in s$. We suggest that this is the critical information to consider when assessing the probability that an observed match is correct. We shall also restrict our attention further to the case when a unique sample unit matches B (so $\tilde{Z}_a = 1$ and $\tilde{Z}_i = 0$ if $i \neq a$ for some unit $a \in s$). This is the worst case and thus of most interest, i.e. the risk will be lower if B matches more than one sample unit. In this case, we obtain the following expression for the identification risk:

$$(2.3) \quad \begin{aligned} \text{Identification risk} &= Pr(A = B|\tilde{Z}_1, \dots, \tilde{Z}_n) \\ &= Pr(E_B) / \sum_{a \in U} Pr(E_a) \end{aligned}$$

where E_a is the event that population unit a is sampled and its value \tilde{X}_a matches X_B and that no other population unit is both sampled and has a value of \tilde{X} which matches X_B . In order to allow for the effect of unequal probability sampling and the potential use of sampling weights, we suppose that units in the population U are selected independently into the sample

s with inclusion probabilities π_j which may depend on the value $\tilde{X} = j$ for the unit. Writing $X_a = k$ and $X_B = j$ and using our previous assumptions about the misclassification mechanism, we obtain $Pr(E_a) = \alpha_j M_{jk} / (1 - \pi_j M_{jk})$, where $\alpha_j = \pi_j \prod_l (1 - \pi_j M_{jl})^{F_l}$ and F_j is the number of units in the population with $X = j$. Hence

$$(2.4) \quad Pr(A = B | \tilde{Z}_1, \dots, \tilde{Z}_n) = [M_{jj} / (1 - \pi_j M_{jj})] / [\sum_k F_k M_{jk} / (1 - \pi_j M_{jk})]$$

This expression assumes the intruder does not know whether $B \in s$. If this event was known to arise and was included in the conditioning set, (2.4) should be modified by setting $\pi_j = 1$ and replacing F_k by f_k . This produces an expression that is similar to that given earlier in (2.2) from Gouweleeuw et al. (1998) but makes fewer approximations. For expression (2.4), the identification risk also assumes that the F_k are part of the data, i.e. known. In practice this will often not be the case, as discussed by Skinner and Shlomo (2008), and it will be necessary to integrate the F_k out of this expression as will be discussed in section 3. It follows from (2.4) that

$$Pr(A = B | \tilde{Z}_1, \dots, \tilde{Z}_n) \leq 1/F_j$$

with equality holding if there is no misclassification. The extent to which the left hand side of this inequality is less than the right hand side measures the impact of misclassification on disclosure risk.

If the inclusion probabilities π_j are all small we may approximate (2.4) by:

$$Pr(A = B | \tilde{Z}_1, \dots, \tilde{Z}_n) = M_{jj} / (\sum_k F_k M_{jk}).$$

Moreover, if the population size is large, we have approximately $\sum_k F_k M_{jk} \approx \tilde{F}_j$, where \tilde{F}_j is the number of units in the population which would have $\tilde{X} = j$ if they were included in the microdata (with misclassification). Hence a simple approximate expression for the risk, natural for many social surveys, is

$$(2.5) \quad Pr(A = B | \tilde{Z}_1, \dots, \tilde{Z}_n) = M_{jj} / \tilde{F}_j.$$

An alternative derivation of this result is provided in the Appendix under the assumption that the intruder adopts the probabilistic record linkage approach of Fellegi and Sunter (1969), making a link if the match variable

$\tilde{Z}_a = 1$. The identification risk is the probability that the match is correct and the above approximation is obtained if the probability is defined with respect to the sampling scheme, the misclassification mechanism and a random selection of a pair for matching as in Fellegi and Sunter (1969).

Another approximation to expression (2.4) is obtained by assuming the misclassification is small, say $M_{jj} = (1 - \delta)\phi_{jj}$ and $M_{jk} = \delta\phi_{jk}$ ($j \neq k$), where the ϕ are fixed and $\delta \rightarrow 0$. In this case, we have

$$(2.6) \quad Pr(A = B | \tilde{Z}_1, \dots, \tilde{Z}_n) \approx F_{jj}^{-1}(1 - [\tilde{F}_j - F_j M_{jj}] / [F_j M_{jj} / (1 - \pi_j M_{jj})])$$

or

$$(2.7) \quad Pr(A = B | \tilde{Z}_1, \dots, \tilde{Z}_n) \approx [M_{jj} / (1 - \pi_j M_{jj})] / [(F_j \pi_j M_{jj}^2) / (1 - \pi_j M_{jj}) + \tilde{F}_j]$$

Note that none of approximations (2.5), (2.6) or (2.7) depend upon M_{jk} for $j \neq k$ and so knowledge of these probabilities is not required in the estimation of risk.

The definition of risk in (2.3) applies to a specific record. Agencies will also usually wish to consider aggregate measures to enable them to make judgements about the whole file. Following Skinner and Shlomo (2008), we define an aggregate measure as the sum of the record-level measures in (2.4) across sample unique records:

$$(2.8) \quad \tau = \sum_{j \in SU} [M_{jj} / (1 - \pi_j M_{jj})] / [\sum_k F_k M_{jk} / (1 - \pi_j M_{jk})],$$

where SU is the set of key variable values which are sample unique. This measure may be interpreted as the expected number of correct matches among sample uniques. For some purposes, an agency might find this measure easier to interpret if it is transformed into a measure with an upper bound, such as by dividing by the number of sample uniques to obtain a proportion. However, we shall stick with the untransformed τ as a measure of the total number of units, e.g. individuals, threatened with identification.

We also consider, for comparison, a related measure which could be used if the misclassification status of microdata records is known. Let $SUCC$ denote the set of key variable values which are sample unique and where these sample unique values have been correctly classified. The measure is given by

$$(2.9) \quad \tau_{CC}^* = \sum_{j \in SUCC} 1/F_j,$$

and again may be interpreted as the expected number of correct matches among sample uniques. We also define τ^* as the corresponding measure of risk in the absence of perturbation, that is the sum of $1/F_j$ across key values which are unique in the sample with respect to X .

3. Risk Estimation

An agency wishing to apply an SDC method to survey microdata will generally not know the values of F_j or \tilde{F}_j appearing in the risk expressions. We do suppose that the values of M_{jk} are known. Skinner and Shlomo (2008) discuss the estimation of risk in the absence of misclassification based on a Poisson - log linear model. In this case, expression (2.4) reduces to $1/F_j$ and their broad approach is to define the risk as the conditional expectation of this quantity given the observed data and to estimate this expectation using data for the sample counts f_j , $j = 1, 2, \dots, K$ for which a log-linear model is fitted. Expression (2.5) provides a simple way to extend their approach to misclassification provided M_{jj} is known. Since the \tilde{f}_j , $j = 1, 2, \dots, K$ represent the available data, all that is required is to ignore the misclassification and estimate the expectation of $1/\tilde{F}_j$ given the data from the \tilde{f}_j , $j = 1, 2, \dots, K$ as in Skinner and Shlomo (2008), that is by fitting a log-linear model now to the \tilde{f}_j , $j = 1, 2, \dots, K$ following the same criteria as before. This results in an estimate $\hat{E}(1/\tilde{F}_j | \tilde{f}_j = 1)$ based on the assumptions of the Poisson distribution for the population and sample counts. These estimates should be multiplied by the M_{jj} values and summed if aggregate measures of the form in (2.8) are needed. It would appear to be rather more complex to estimate the expressions including terms in F_j . In the presence of complex sampling, the estimation method may be adapted using the method of pseudo maximum likelihood estimation (Rao and Thomas, 2003) by incorporating survey weights in the estimation as discussed by Skinner and Shlomo (2008).

4. Application of Perturbative Disclosure Limitation Techniques

In this section we consider two specific perturbative SDL techniques used at statistical agencies: data swapping and the post-randomization method (PRAM). Both techniques introduce misclassification of the key variables to lower the probabilities of identifying individuals. We present examples of how to assess the impact of these techniques on identification risk. Since the misclassification is under the control of the statistical agency, the misclassification matrix M is known.

4.1. Data swapping

The method of data swapping is based on exchanging the values of one or more key variables between pairs of records. In order to minimize bias, the pairs of records are typically selected within strata defined by control variables, such as the age and sex of the individual. In addition, the perturbation can be targeted to high-risk records that are more likely to be population uniques, for example on rare ethnicities. It is common that geographic variables are swapped between records for the following reasons:

- For given values of the control variables, the sensitive variables are likely to be relatively independent of geography and therefore it is expected that less bias will occur. In addition, swapping geography will not normally result in inconsistent and illogical records. By contrast, swapping a variable such as age would result in many inconsistencies with other variables, such as marital status and education.
- At a higher geographical level and within control strata, the marginal distributions are preserved.
- The level of protection increases by swapping variables which are highly 'matchable' such as geography.

For this experiment, we carry out a simple data swapping procedure where the geography variable of Local Authority District (LAD) is exchanged between a pair of individuals. The population includes $N = 1,468,255$ individuals from an extract of the 2001 United Kingdom (UK) Census. We drew 1% Bernoulli samples ($n = 14,683$) and define six key variables for the risk assessment: Local Authority (LAD) (11), sex (2), age groups (24), marital status (6), ethnicity (17), economic activity (10), where the numbers of categories of each variable are in parentheses ($K = 538,560$). We implement a random data swap by drawing a sub-sample of 10% and 20% in each of the LADs. The remaining individuals are not perturbed. On the sub-samples in each LAD, half of the individuals are flagged. For each flagged individual, an unflagged individual is randomly chosen within the sub-sample and their LAD variables swapped, on condition that the individual chosen was not previously selected for swapping and that the two individuals do not have the same LAD, i.e. no individual is selected twice for producing a pair. We also implemented a 10% and 20% targeted data swap where the LAD variable is swapped separately within two groups defined by 'White British' and 'Other' ethnicities. For the 20% swap, LADs were swapped randomly between all pairs of individuals in the 'Other' group and a small percentage (7%) of individuals in the 'White British' group. This swapping rate was chosen so that the total percentage of swapped individuals would be 20%

as in the random data swapping. For the 10% swap, LADs were swapped randomly from among the 'Other' group that compose 10% of the total individuals in the sample.

The misclassification matrix M for the data swapping designs can be expressed simply in terms of the 11×11 misclassification matrix, denoted $M^g = [M_{jk}^g]$, for the LAD variable g :

For the random swap:

- On the diagonal: $M_{jj}^g = 0.9$ or $M_{jj}^g = 0.8$ for the 10% and 20% swaps respectively.
- Off the diagonal: $M_{jk}^g = 0.1 \times n_k / (\sum_{l \neq j} n_l)$ or $M_{jk}^g = 0.2 \times n_k / (\sum_{l \neq j} n_l)$, where n_k is the number of records in the sample in LAD k , $k = 1, 2, \dots, 11$, for the 10% and 20% swaps respectively.

For the targeted swap on the 10% swap, the values M_{jk}^g for the 'Other' ethnicity are calculated as follows:

- On the diagonal: $M_{jj}^g = .25$
- Off the diagonal: $M_{jk}^g = .75 \times n_{2k} / (\sum_{l \neq j} n_{2l})$, where n_{2k} is the number of records in the sample with 'Other' ethnicity in LAD k , $k = 1, 2, \dots, 11$.

For the targeted swap on the 20% swap, the misclassification matrix M is defined separately according to the 'White British' and 'Other' ethnicities as follows:

- On the diagonal: $M_{jj}^g = 0.93$
- Off the diagonal: $M_{jk}^g = 0.07 \times n_{1k} / (\sum_{l \neq j} n_{1l})$, where n_{1k} is the number of records in the sample with 'White British' ethnicity in LAD k , $k = 1, 2, \dots, 11$.

The values M_{jk}^g for the 'Other' ethnicity are calculated as follows:

- On the diagonal: $M_{jj}^g = 0$
- Off the diagonal: $M_{jk}^g = 1 \times n_{2k} / (\sum_{l \neq j} n_{2l})$, where n_{2k} is the number of records in the sample with 'Other' ethnicity in LAD k , $k = 1, 2, \dots, 11$.

4.2. The Post-Randomization Method (PRAM)

A more direct method that is used for exchanging values of categorical variables is PRAM. For this method, values of categories in a given record are changed or not changed stochastically according to a misclassification matrix. This matrix is chosen to preserve expected marginal frequencies of the variables. Let f^c be the row vector of sample frequencies of the different

categories of key variable X^c and $p^c = f^c/n$ be the corresponding vector of sample proportions, where n is the sample size. For each record, the category of X^c is changed or not changed according to the probabilities in the misclassification matrix M^c . Let \tilde{f}^c be the row vector of perturbed frequencies. Then $E(\tilde{f}^c|f^c) = f^c M^c$, where the expectation is with respect to the misclassification mechanism. Provided that the matrix M^c has an inverse $M^{c^{-1}}$, this can be used to obtain an unbiased estimator of the original frequency vector: $\hat{f}^c = \tilde{f}^c M^{c^{-1}}$. In addition, we can place the condition of invariance on the matrix M^c , i.e. $f^c M^c = f^c$, and preserve the expected marginal frequencies. This releases the users of the perturbed file of the extra effort to obtain unbiased moment estimates of the original data, since \tilde{f}^c itself will be an unbiased estimate of f^c .

To obtain an invariant transition matrix, the following two stage algorithm is applied (see Willenborg and De Waal, 2001). Let M^c be the misclassification matrix: $M_{jk}^c = Pr(\tilde{X}^c = k|X^c = j)$ where j represents the original category and k the perturbed category. Now calculate the matrix Q using Bayes formula by $Q_{kj}^c = Pr(X^c = j|\tilde{X}^c = k) = M_{jk}^c Pr(X^c = j)/[\sum_l M_{lk}^c Pr(X^c = l)]$. We estimate the entries of this matrix by $\hat{Q}_{kj}^c = M_{jk}^c p_j^c / [\sum_l M_{lk}^c p_l^c]$, where p_j^c is the sample proportion in category j . The matrix $R^c = M^c \hat{Q}^c$ is invariant, i.e. $p^c R^c = p^c$, since $R_{ij}^c = \sum_k [p_j^c M_{ik}^c M_{jk}^c / \sum_l M_{lk}^c p_l^c]$ and $\sum_i p_i^c R_{ij}^c = \sum_k p_j^c M_{ik}^c = p_j^c$. The vector of the original proportions p^c is the eigenvector of R . In practice, \hat{Q}^c can be calculated by transposing matrix M^c , multiplying each column j by p_j^c and then normalizing its rows so that the sum of each row equals one. We define $R^{c*} = \alpha R^c + (1 - \alpha)I$ where I is the identity matrix of the appropriate size. R^{c*} is also invariant and the amount of misclassification is controlled by the value of α .

We conduct a second experiment using the same data and set-up described in Section 4.1 and PRAM to perturb the geographical variable LAD. For the random perturbation, an 11×11 misclassification matrix M^c is defined for the 11 categories of LAD where the diagonal elements are 0.9 and 0.8 and the off-diagonal elements are equal to a probability of 0.1 and 0.2 for the 10% and 20% perturbation respectively. The invariant misclassification matrix is calculated with $\alpha = 0.55$. For each individual, a random uniform number between 0 and 1 is generated and the category of the LAD changed (or not changed) if it is in the interval defined by the cumulative probability. For the 10% targeted perturbation, we define the misclassification matrix for the 'Other' ethnicities with 0.25 on the diagonal and 0.75 on the off-diagonals and the invariant parameter $\alpha = 0.85$. For the 20% targeted perturbation, we define the misclassification matrix for the 'Other' ethnicities with 0 on

the diagonal and 1 on the off-diagonals, and the misclassification matrix for the 'White British' ethnicity with 0.93 on the diagonal and 0.07 on the off-diagonals. For both matrices, the invariant parameter is $\alpha = 1$.

4.3. Results of Disclosure Risk Assessment

Since we know the misclassification matrix M and the true population counts F_j in these experiments, we can assess the performance of expressions (2.5)-(2.7) as approximations to (2.4). We do this by summing all the expressions across sample unique records, as in the aggregate risk measure τ in (2.8) and comparing the resulting sums. We also compare these measures to the measure in (2.2) of Gouweleeuw et al. (1998). In addition, we consider the more practical situation when neither the F_j nor the \tilde{F}_j are known to the agency, all that is observed is the 'misclassified' sample and the matrix M . In this case, we carry out the risk estimation as described in Section 3 through the use of the Poisson-log-linear model on the sample counts \tilde{f}_j . The log-linear model was chosen using a forward search algorithm and the outcome of goodness of fit statistics as developed in Skinner and Shlomo (2008). We calculate the naive estimated risk measure obtained from the log-linear model on the misclassified sample and the adjusted estimated risk measure taking into account the misclassification matrix. The experiments were repeated under different samples and each perturbation method applied independently and we found that all of the experiments produced similar results. Table 1 presents results of one of the simulation experiments for each of the perturbation methods: random and targeted data swapping and PRAM.

The estimates presented in Table 1 for the risk of identification are similar for random data swapping and PRAM. Misclassification reduces the risk in the file from about $\tau^* = 360$ to about $\tau_{CC}^* = 290$ for the 20% perturbation and $\tau_{CC}^* = 320$ for the 10% perturbation for those methods. The measure τ^* is interpreted as the expected number of correct matches which an intruder would make if matches were attempted with all sample unique records. The decrease in this measure from 360 to 290 as a result of misclassification is modest since a large number of records remain unchanged. An alternative interpretation of τ could be obtained by dividing by the number of sample uniques to give the proportion of sample uniques which would be expected to be identified correctly. This proportion ranges in Table 1 between 0.116 for the 10% Random Swap, 0.053 for the 10% Targeted Swap, 0.108 for the 20% Random Swap and 0.030 for the 20% Targeted Swap.

The identification risk is reduced considerably with the targeted data

TABLE 1

Identification risk estimates for microdata samples generated from UK 2001 Census subject to perturbative SDL methods - Risk measure τ^* no misclassification = 358.1

Identification Risk Measures	SDL Method			
	Random		Targeted	
	Swap	PRAM	Swap	PRAM
10% Perturbation				
Identification risk measures for perturbed data with known population counts				
<i>Risk measure τ in (2.8)</i>	321.6	325.8	146.3	161.6
<i>Approximation in (2.5)</i>	321.4	325.5	146.2	161.4
<i>Approximation in (2.6)</i>	317.7	321.7	144.8	159.8
<i>Approximation in (2.7)</i>	321.6	325.6	146.3	161.6
<i>Risk measure τ_{CC}^* in (2.9)</i>	316.6	318.2	149.5	160.3
Estimated risk measures based on sample data				
<i>Risk measure in (2.2)</i>	2,486.7	2,489.1	1,749.1	1,899.3
<i>Naive risk measure (Poisson – log linear model on misclassified sample)</i>	343.2	347.6	297.2	285.4
<i>Estimated risk measure (Poisson – log linear model adjusted for misclassification)</i>	308.8	312.7	142.7	157.9
20% Perturbation				
Identification risk measures for perturbed data with known population counts				
<i>Risk measure τ in (2.8)</i>	298.9	299.7	82.2	133.8
<i>Approximation in (2.5)</i>	298.4	299.3	82.1	133.7
<i>Approximation in (2.6)</i>	280.4	283.5	81.7	132.7
<i>Approximation in (2.7)</i>	298.9	299.8	82.2	133.8
<i>Risk measure τ_{CC}^* in (2.9)</i>	292.8	292.2	85.0	133.4
Estimated risk measures based on sample data				
<i>Risk measure in (2.2)</i>	2,264.0	2,311.7	1,419.8	1,688.2
<i>Naive risk measure (Poisson – log linear model on misclassified sample)</i>	358.6	349.5	262.5	285.2
<i>Estimated risk measure (Poisson – log linear model adjusted for misclassification)</i>	286.8	283.1	90.3	133.2

TABLE 2
Cross-classification of per-record risk measures in (2.4) against estimates based on Poisson log-linear model adjusted for misclassification for sample uniques under 20% random data swap

Per-record risk Measures from 2.4	Estimates from Poisson Log-linear Model			
	0.00 – 0.09	0.10 – 0.49	0.50 – 1.00	Total
0.00 – 0.09	1,961	133	4	2,098
0.10 – 0.49	180	325	76	581
0.50 – 1.00	8	69	75	152
Total	2,149	527	155	2,831

swapping since many more sample uniques are perturbed. The misclassification is reduced from about $\tau^* = 360$ to about $\tau_{CC}^* = 85$ for data swapping and $\tau_{CC}^* = 130$ for PRAM for the 20% perturbation and to about $\tau_{CC}^* = 150$ for data swapping and $\tau_{CC}^* = 160$ for PRAM for the 10% perturbation. The three approximations to the risk measure in (2.8) all provide good results although the approximation in (2.6) slightly underestimates. The measure in (2.8) relies on knowledge of both the full misclassification matrix M and the population counts F_j . In contrast, the approximations (2.5), (2.6) and (2.7) only require knowledge of the probability of not misclassifying a record, i.e. the probabilities on the diagonals. The alternative risk measure τ_{CC}^* in (2.9) also turns out to behave similarly to (2.8). The value of the measure in (2.2) of Gouweleeuw et al. (1998) is much higher than the values of the other measures, reflecting the very conservative assumption that the intruder knows that the target unit is in the microdata sample. In practice the population counts will generally be unknown to the statistical agency (and the intruder) for survey data. We therefore consider the method in Section 3 based upon the Poisson-log linear model. The estimated aggregate risk measures appear to perform well with estimates for the risk measure under misclassification of about 285 for random data swapping and PRAM under the 20% perturbation and about 310 for random data swapping and PRAM under the 10% perturbation. The estimated aggregate risk measures are about 140 for targeted data swapping and 160 for targeted PRAM for the 20% perturbation and about 90 for targeted data swapping and 130 for targeted PRAM for the 10% perturbation.

Another important consideration when assessing disclosure risk for releasing microdata is the individual per-record (record-level) disclosure risk measures in (2.4). Individual records with high disclosure risk might be subjected to further tailored perturbation. In Figure 1, we plot the per-record (record-level) risk measures in (2.4) for the sample uniques against the estimated adjusted risk measures (as described in Section 3) based on the Poisson-log-

linear model for the experiment based on 20% random data swapping. In addition, we summarize this bivariate distribution for the sample uniques in a two-way table in Table 2. In both of the analyses we see a good fit between the risk measures in (2.4) and their estimated risk measures. The Spearman's rank correlation is 0.91.




FIG 1. Scatterplot of individual per-record risk measures in (2.4) against estimated risk measures based on Poisson-log linear model under 20% random data swap

4.4. Results of Information Loss Assessment

The utility of microdata that has undergone data masking techniques is measured here in terms of the closeness of the results of an analysis based upon the perturbed data compared to the same analysis based upon the original data. The nature of the results and the type of analysis depend on user requirements. In general, microdata is multi-purpose and used by many different users. For this assessment we use the following three information loss measures reflecting distortions of distributions in two-way tables, as considered by Gomatam and Karr (2003) and Shlomo and Young (2006):

- Relative Absolute Average Distance per Cell: Let D represent a frequency distribution for a two-way table produced from the microdata and let $D(r, c)$ be the frequency in the cell in row r and column c . The distance metric is:

$$RAAD(D_{orig}, D_{pert}) = 100 \times (D_{avg} - AAD) / D_{avg}$$

where the average cell size is defined as:

$$D_{avg} = \sum_{r,c} D_{orig}(r,c)/RC$$

with R the number of rows and C the number of columns in the table, and the AAD metric is defined as:

$$AAD(D_{orig}, D_{pert}) = \sum_{r,c} |D_{pert}(r,c) - D_{orig}(r,c)|/RC$$

with *pert* and *orig* referring to the perturbed and original tables respectively. The *RAAD* provides a measure of the average absolute perturbation per cell compared to the average cell size of the table.

- Impact on Measures of Association:

$$RCV(D_{orig}, D_{pert}) = 100 \times (CV(D_{pert}) - CV(D_{orig}))/CV(D_{orig})$$

where

$$CV(D) = \sqrt{\chi^2/\min(R-1, C-1)}$$

is Cramer's measure of association, defined in terms of χ^2 , the usual Pearson chi-squared statistic for testing independence in the two-way table. The *RCV* provides a measure of attenuation of the association.

- Impact on an ANOVA Analysis: another form of bivariate analysis consists of comparing proportions in a category of a column (outcome) variable between categories of a row (explanatory) variable. Let $P^c(r) = D(r,c)/\sum_c D(r,c)$ be the proportion in column c for row r and define the between-row variance of this proportion by:

$$BV(P^c) = \sum_r (P^c(r) - P^c)^2/(R-1)$$

where $P^c = \sum_r D(r,c)/\sum_{rc} D(r,c)$. The measure of information loss is:

$$BVR(P_{orig}^c, P_{pert}^c) = 100 \times (BV(P_{pert}^c) - BV(P_{orig}^c))/BV(P_{orig}^c)$$

The *BVR* provides a measure of attenuation of between group differences in an ANOVA analysis.

Table 3 presents results of the information loss measures on the misclassified samples used in Table 1. We obtain similar results for the information loss measures on the random perturbations between data swapping and PRAM with an expected improvement under the smaller perturbation

TABLE 3
 Information loss measures for microdata samples generated from UK 2001 Census
 subject to three perturbative SDL methods

Information Loss Measures	SDL Method			
	Random		Targeted	
	Swap	PRAM	Swap	PRAM
10% Perturbation				
<i>RAAD on LAD × ethnicity</i>	98.5	98.1	97.4	97.2
<i>RAAD on LAD × economic activity</i>	97.0	96.9	96.1	95.8
<i>RCV on LAD × ethnicity</i>	-9.9	-10.4	-13.3	-12.9
<i>RCV on LAD × economic activity</i>	-10.8	-9.8	-11.0	-10.4
<i>BVR on prop. 'White British' across LAD</i>	-20.9	-23.8	0	0
<i>BVR on prop. 'Indian' across LAD</i>	-12.6	-13.0	-18.9	-17.3
20% Perturbation				
<i>RAAD on LAD × ethnicity</i>	97.4	97.2	96.5	96.4
<i>RAAD on LAD × economic activity</i>	95.8	95.5	95.0	94.9
<i>RCV on LAD × ethnicity</i>	-20.4	-20.4	-17.8	-16.9
<i>RCV on LAD × economic activity</i>	-18.1	-17.0	-16.2	-14.4
<i>BVR on proportion 'White British' across LAD</i>	-37.4	-39.6	0	0
<i>BVR on proportion 'Indian' across LAD</i>	-37.5	-39.1	-34.2	-29.5

rate of 10%. The targeted perturbation show slight improvements on the *RCV* and *BVR* for the 20% perturbation but there is generally no improvement under the 10% perturbation. The impact on the *BVR* for other ethnic groups was mixed with most of the ethnic groups following the same pattern of attenuation as seen for the 'Indian' ethnic group under data swapping and PRAM but with a few exceptions due to small sample sizes. For example, we obtained a positive value of the *BVR* for the 'Chinese' ethnicity. Overall, the considerable reduction in disclosure risk achieved by the 20% targeted data swapping in Table 1 does not appear to be offset by any major reduction in utility compared to the other methods.

In Figure 2, we plot a risk-utility map (Duncan, et al., 2001). The points on the map represent different candidate releases, i.e. perturbation methods with different levels of perturbation. In addition to the levels considered earlier, we also include 2% and 5% targeted and random perturbation. The points are denoted *T* for targeted or *R* for random; 20 for 20%, 10 for 10%, 5 for 5% or 2 for 2%; and *S* for swapping or *P* for PRAM. The points are plotted against the risk measure τ in (2.8) on the Y-Axis and the information loss measure *RAAD* for *LAD × ethnicity* on the X-Axis. We see that at the same level of information loss between the targeted 10% perturbation and the random 20% perturbation with respect to the *RAAD*, we obtain lower disclosure risk with the targeted 10% perturbation. The same applies to the targeted 5% perturbation and the random 10% perturbation with

the targeted 5% perturbation having less disclosure risk than the random 10% perturbation at the same level of information loss. We draw a line to connect points on the risk-utility frontier (Gomatam, et al., 2005) and note that in all cases, at given levels of information loss, the targeted data swapping provides the lowest disclosure risk compared to the other methods, although there is little difference between targeted swapping and targeted PRAM. Targeting did not appear to lead to much greater information loss for the other measures in Table 3 and the general conclusion here is that targeting seems useful, enabling less perturbation to be applied and hence less information loss for a given level of risk protection. Of course, this finding could vary in other settings and an agency could use a similar risk-utility approach, based on its own data, to determine its preferred SDL approach.




FIG 2. Risk-Utility Map

5. Impact of Misclassifying Multiple Key Variables

The previous section only provided estimates of the impact of misclassifying one key variable. In this section we provide a further numerical illustration to demonstrate the potential impact of misclassifying multiple key variables. We consider a simple set-up where the C key variables X^1, \dots, X^C are independent and binary. Their values in the external information and the microdata are denoted X^c and \tilde{X}^c respectively, $c = 1, \dots, C$. We suppose that $Pr(X^c = 2) = p$, $Pr(X^c = 1) = 1 - p$, $Pr(\tilde{X}^c = 2|X^c = 1) = \theta_1$ and

$Pr(\tilde{X}^c = 1 | X^c = 2) = \theta_2$ for $c = 1, \dots, C$. The misclassification probabilities M_{jk} in (2.1) will thus consist of products of C terms, each term being one of $\theta_1, 1 - \theta_1, \theta_2$ or $1 - \theta_2$. To force X and \tilde{X} to share the same marginal distribution we set $\theta_2 = (1-p)\theta_1/p$ so that $Pr(\tilde{X}^c = 1) = p$ and, to simplify, write $\theta_1 = \theta$.

In our experiment we generated values of X for a population of size N , drew a sample of size n by simple random sampling and then generated the values \tilde{X} . Various choices of (N, n, C, p, θ) were considered. We also generated \tilde{X} for all population units so that \tilde{F}_j could be computed.




FIG 3. Risk measure for different numbers of key variables and rates of misclassification

We report values of risk measure (2.5) summed over sample uniques $\sum_{SU} M_{jj}/\tilde{F}_{jj}$ in Figure 3 for $N = 100,000$, $n = 2,000$, $p = 0.2$ and for various choices of C and θ . Note that the number of sample uniques increases as we add in more binary key variables. For $C = 11$ we have about 240 sample uniques and for $C = 30$ we have about 1,960 sample uniques. In the absence of misclassification we find that the risk increases monotonically and rapidly with C . This is because the number of population uniques is increasing with C and the fact that any observed match with a population unique must be a true match. On the other hand, in the presence of misclassification, we find that the risk does not increase monotonically, rather it reaches a maximum and then declines. As expected, the more misclassification, the lower the disclosure risk.

We do not present information loss measures for the simulation since their

values follow theoretically. For any analysis involving a given set of variables, say the estimation of a table cross-classifying two particular key variables, the addition of further key variables will have no systematic impact on any of the information loss measures, since each of the variables of interest will be perturbed in the same way, irrespective of the inclusion of other key variables. The only variation we might expect to observe would be as a result of simulation variation. Any information loss function in Figure 3 should therefore be flat.

6. Discussion

In this paper, we have shown how existing methods for assessing identification risk in survey microdata may be extended in a relatively simple way to capture the impact of SDL methods based on misclassification. We presented a general expression for the risk under misclassification in (2.4) and showed that the simple formula in (2.5) provided a good approximation to this expression in two experiments based upon UK census data. The advantage of the formula in (2.5) is that it enables the extension of existing risk assessment methods for unperturbed data based on Poisson log-linear models, as discussed in Skinner and Shlomo (2008), to handle perturbative SDL methods. We demonstrated this extended approach also with the census data and provided a disclosure risk-data utility analysis. We showed how a targeted SDL method could dominate corresponding random methods.

One challenge faced by agencies when assessing identification risk is the need to make assumptions about the information available to the intruder, specifically the nature and number of key variables. We conducted a numerical experiment to assess the sensitivity of the identification risk to the misclassification of different numbers of key variables. In the absence of misclassification, the risk can increase rapidly with the number of key variables. We observed that misclassification can, however, dominate this effect with the risk eventually declining as the number of key variables increases. This is potentially an encouraging finding for agencies, since the sensitivity of the identification risk to departures from assumptions about the choice of key variables may be reduced in some settings when the kinds of SDL methods considered here are used and, in cases such as in Figure 3, there may even be a natural upper bound for the risk across plausible choices.

Another issue faced by agencies is whether to release values of the parameters of the SDL method employed, e.g. the swapping rate. The information loss measures used in section 4.4 assume that users of the microdata simply ignore the perturbation in their analyses of the data. The agency's aim is to

find an SDL method for which both the information loss and the disclosure risk are considered satisfactorily small. If this is not feasible then it may be necessary for the agency to resort to an SDL method which leads to non-negligible distortion of analyses. In this case it may be desirable for data analysts to be provided with values of the parameters of the SDL method to enable them to undertake valid inference, as discussed e.g. in Gouweleeuw et al. (1998) for PRAM (note that our use of invariant PRAM was designed to avoid this need). The disclosure risk implications of releasing such SDL parameters will not be pursued further here.

The findings of this paper are not only relevant to understanding the impact of SDL methods but also to the assessment of risk, before the application of SDL methods, in a way which more realistically takes account of the errors of classification which arise in survey data from measurement, coding and processing as well as from imputation for missing data, providing the agency has estimates for the diagonal elements of the misclassification matrix.

Appendix: Derivation of (2.5) under Probabilistic Record Linkage

Suppose, as before, that a microdata record i is linked to a target unit B by comparing the values of \tilde{X}_i and X_B . Following the approach of Fellegi and Sunter (1969), let $\gamma(\tilde{X}_i, X_B) = j$ if $\tilde{X}_i = X_B = j$, $j = 1, \dots, K$ and $\gamma(\tilde{X}_i, X_B) = K + 1$ if $\tilde{X}_i \neq X_B$ and suppose that exact matching is used, so that a link is made if $\gamma(\tilde{X}_i, X_B) \leq K$. Suppose the intruder draws the pair (i, B) at random (with equal probability) from the set of pairs $s \times s^*$, where s^* is the subset of units in U appearing in the external database from which the intruder selects B . Partition $s \times s^*$ as $M = \{(i, B) | i = B\}$ and $U = \{(i, B) | i \neq B\}$ and let $m(j) = \Pr[\gamma(\tilde{X}_i, X_B) = j | (i, B) \in M]$, $u(j) = \Pr[\gamma(\tilde{X}_i, X_B) = j | (i, B) \in U]$ and $p = \Pr[(i, B) \in M]$, where $\Pr(\cdot)$ is defined with respect to the selection of (i, B) , the selection of the sample s and the misclassification mechanism. Then the identification risk for a linked pair (i, B) for which $\tilde{X}_i = X_B = j$ is given by

$$\phi_j = \Pr[(i, B) \in M | (\tilde{X}_i, X_B) = j] = \frac{m(j)p}{m(j)p + u(j)(1-p)}$$

A large sample size approximation gives $m(j) \approx M_{jj}f_j/n^*$, $u(j) \approx (\pi\tilde{F}_j f_j - \pi M_{jj} f_j)/(nn^* - \pi n^*)$, $p = \pi/n$, where f_j is the number of units b in s^* for which $X_b = j$ and n^* is the size of s^* . It follows that $\phi_j \approx M_{jj}/\tilde{F}_j$ irrespective of the manner in which s^* is selected from U . Skinner (2008) provides

further discussion of identification risk under probabilistic record linkage.

References

- [1] DALENIUS, T. and REISS, S. P. (1982). *Data Swapping: a Technique for Disclosure Control*. Journal of Statistical Planning and Inference, **6**, pp 73–85.
- [2] DOMINGO-FERRER, J. and TORRA, V. (2001). *A Quantitative Comparison of Disclosure Control Methods for Microdata* in (P. Doyle, J. Lane, J. Theeuwes, and L. Zayatz, eds.) Confidentiality, Disclosure Control and Data Access: Theory and Practical Applications, Amsterdam, The Netherlands: North Holland, pp 111–145.
- [3] DUNCAN, G. and LAMBERT, D. (1989). *The Risk of Disclosure for Microdata*. Journal of Business and Economic Statistics, **7**, pp 207–217.
- [4] DUNCAN, G., KELLER-MCNULTY, S. and STOKES, S. (2001). *Disclosure Risk vs. Data Utility: the R-U Confidentiality Map*. Technical Report LA-UR-01-6428, Statistical Sciences Group, Los Alamos, N.M.: Los Alamos National Laboratory.
- [5] ELAMIR, E. and SKINNER, C. (2006). *Record-Level Measures of Disclosure Risk for Survey Micro-data*. Journal of Official Statistics, **22**, pp 525–539.
- [6] FELLEGI, I. and SUNTER, A. (1969). *A Theory for Record Linkage*. Journal of the American Statistical Association, **64**, pp 1183–1210.
- [7] FIENBERG, S.E. (2006). *Privacy and confidentiality in an e-commerce world: data mining, data warehousing, matching and disclosure limitation* Statistical Science, **21**, pp 143–154.
- [8] FULLER, W. (1993). *Masking Procedures for Micro-data Disclosure Limitation*. Journal of Official Statistics, **9**, pp 383–406.
- [9] GOMATAM, S. AND KARR, A. (2003). *Distortion Measures for Categorical Data Swapping*. Technical Report Number 131, National Institute of Statistical Sciences.
- [10] GOMATAM, S., KARR, A. and SANIL, A.P. (2005). *Data Swapping as a Decision Problem*. Journal of Official Statistics, **21**, pp 635–655.
- [11] GOUWELEEUW, J., KOOIMAN, P., WILLENBORG, L.C.R.J. and DE WOLF, P.P. (1998). *Post Randomisation for Statistical Disclosure Control*. Journal of Official Statistics, **14**, pp 463–478.
- [12] KUHA, J., and SKINNER, C. (1997). *Categorical Data Analysis and Misclassification*. in (L. Lyberg, P. Biemer, M. Collins, E. De Leeuw, C. Dippo, N. Schwarz, and D. Trewin, eds.) Survey Measurement and Process Quality, New York: Wiley, pp 633–670.

- [13] NATIONAL RESEARCH COUNCIL(2005). *Expanding Access to Research Data: Reconciling Risks and Opportunities*. Washington, DC; The National Academies Press.
- [14] PAASS, G.,(1988). *Disclosure Risk and Disclosure Avoidance for Microdata*. Journal of Business and Economic Statistics, **6**, pp 487–500.
- [15] RAO, J.N.K AND THOMAS, D.R. (2003). *Analysis of Categorical Response Data from Complex Surveys: an Appraisal and Update* in (Chambers, R.L. and Skinner, C.J. eds.) *Analysis of Survey Data*, UK: Wiley and sons, pp 85–108.
- [16] REITER, J. (2005). *Estimating Risks of Identification Disclosure in Microdata*. Journal of the American Statistical Association, **100**, pp 1103–1112.
- [17] SHLOMO, N. and YOUNG, C. (2006). *Statistical Disclosure Limitation Methods Through a Risk-Utility Framework* in (Domingo-Ferrer, J. and Franconi, L. eds.) *Privacy in Statistical Databases*, New York: Springer, LNCS **4302**, pp 68–81.
- [18] SKINNER, C.J. (2008). *Assessing Disclosure Risk for Record Linkage* in (Domingo-Ferrer, J., ed.) *Privacy in Statistical Databases 2008 Lecture Notes in Computer Science*, New York: Springer-Verlag.
- [19] SKINNER, C. and HOLMES, D. (1998). *Estimating the Re-identification Risk per Record in Microdata*. Journal of Official Statistics, **14**, pp 361–372.
- [20] SKINNER, C. and SHLOMO, N. (2008). *Assessing Identification Risk in Survey Microdata Using Log-linear Models*. Journal of the American Statistical Association, **103**, pp 989–1001.
- [21] SPRUILL, N.L. (1982). *Measures of Confidentiality*. Proceedings of the Survey Research Methods Section of the American Statistical Association, pp 260–265.
- [22] WILLENBORG, L.C.R.J.. and DE WAAL, T. (2001). *Elements of Statistical Disclosure Control in Practice* Lecture Notes in Statistic 155. New York: Springer-Verlag.
- [23] WINKLER, W.E. (2004). *Re-identification Methods for Masked Microdata* in (Domingo-Ferrer, J. and Torra, V., eds.) *Privacy in Statistical Databases. Lecture Notes in Computer Science 3050*, Berlin: Springer, pp 216–230.
- [24] YANCEY, W.E., WINKLER, W.E. and CREECY, R.H. (2002). *Disclosure Risk Assessment in Perturbation Micro-Data Protection* in (Domingo-Ferrer, J., ed.) *Inference Control in Statistical Databases*, New York: Springer, pp 135–151.

University of Southampton
Southampton, SO17 1BJ United Kingdom
e-mail: C.J.Skinner@Soton.ac.uk
N.Shlomo@Soton.ac.uk