I^2SDS
The Institute for Integrating Statistics in Decision Sciences

Technical Report TR-2008-12
August 29, 2008

Reflections on Data Perturbation and Post-Randomization for Statistical Disclosure Control

Tapan K. Nayak
Department of Statistics
The George Washington University, USA

REFLECTIONS ON DATA PERTURBATION AND POST-RANDOMIZATION FOR STATISTICAL DISCLOSURE CONTROL

Tapan K. Nayak

Department of Statistics

George Washington University

Washington, DC 20052, USA.

Abstract

Protection against disclosure of confidential information is important for statistical agencies releasing microdata from census and sample surveys. We examine the perturbation method and the post-randomization method (PRAM), two recently introduced techniques for disclosure control. We find that the previous claims that the perturbation method preserves all statistical information and that it does not increase disclosure risk are questionable. The post-randomization method is similar to randomized response surveys but the issue of privacy protection is different in the two contexts. We clarify this point and discuss suitable approaches for assessing disclosure risk after using PRAM.

Key words and Phrases: Data utility; Disclosure risk; Randomized response; Respondents' privacy; Unbiased estimation; Variance inflation.

1. Introduction.

The primary objective of most statistical agencies is to collect and publish data relevant to important national and regional public policy issues. Dissemination of statistical data informs the public and policy makers and facilitates research and creation of new knowledge in many areas, helpful for improving social and economic systems, physical infrastructure and public policy. While suitably selected summary statistics and graphs may well inform the public and policy makers, researchers often require microdata to carry out their own investigations. The agencies, however, may not release the original microdata as they also need to protect the privacy of survey respondents for legal reasons and for upholding public trust. Agencies try to provide protection from the risk of disclosure by releasing a modified (or masked) version of the original data, at the expense of some statistical information loss. Thus, the agencies face a fundamental trade-off between privacy protection and the utility of released data.

For privacy protection, obviously all direct identifiers, such as full name, passport number and social security number, must be removed prior to data release. But, that is not sufficient for disclosure avoidance as it may be possible to identify the record of a survey participant by matching the values of variables such as gender, birth date and zip code, that are easily available from other sources. Often the original data need to be further modified for protection against disclosure, and various procedures, such as grouping, collapsing response categories, data swapping, multiple imputation, post-randomization method (PRAM) and random noise addition have been used in practice. The books by Doyle et al. (2001) and Willenborg and de Waal (2001) discuss many important issues relating to disclosure and various disclosure control techniques.

Typically, a microdata set contains records of n individuals (or sampling units) on k variables, some of which are confidential or sensitive. Some attributes such as gender, age and marital status, are known easily from other sources and are useful for identification. Such variables are generally non-sensitive and will be called key variables, following Bethlehem et al. (1990). For simplicity, we shall assume that all survey variables can be divided into two groups: key variables and confidential variables, although in practice, this classification may be subjective and dubious. We shall denote all key variables by X and all confidential variables by Y. The survey variables may be nominal, ordinal or quantitative and the nature of the variables may affect both disclosure risk and data utility. For simplicity and clarity of exposition we shall often consider only qualitative variables.

Disclosure is a difficult topic (cf., Lambert, 1993) and it can occur in different forms associated with different disclosure scenarios (cf., Willenborg and De Wall, 2001). Broadly speaking, disclosure occurs when the values of some confidential variables for a specific subject can be predicted *too accurately* based on the values of the key variables X. The most serious type of disclosure is identity disclosure, which happens when the microdata record of a subject can be discerned from the values of the key variables. Identity disclosure reveals the values of all confidential variables of an identified subject. Identity disclosure is closely related to an unit being unique in the population (or sample) with respect to the key variables (Bethlehem et al., 1990; Greenberg and Zayatz, 1992). If an unit is not unique in the sample, its identity cannot be ascertained with certainty. Measures of identity disclosure center around the proportion of sample or population units whose records can be identified by matching the key variables. Assessment of identity disclosure risk has been discussed by Bethlehem et al. (1990), Greenberg and Zayatz (1992), Willenborg and De Wall (2001), Skinner and Elliot (2002), Reiter (2005)

and others. In particular, Skinner and Elliot (2002) discussed three measures and recommended to use the proportion of correct matches among those population units which match a sample unique microdata record for assessing identity disclosure risk.

Another type of disclosure that has received much attention is predictive disclosure, which occurs if the released data enable one to infer a respondent's value of a confidential variable with high accuracy. An extreme case of predictive disclosure is attribute disclosure, where a confidential variable value can be predicted completely accurately. Attribute disclosure can occur without identity disclosure. For example, suppose in a data set, a confidential variable has the same value, say y_0, for all respondents with a specific value, say x_0, of the key variables. Then, if the original data are released, an intruder would know surely the confidential variable value (y_0) of any respondent with $X = x_0$, but not the identify of the respondent in the data set. The essence of predictive disclosure is: an intruder gaining too much new information about specific respondents from the released data. So, predictive disclosure depends not only on the released data set but also on the intruder's prior knowledge, and hence predictive disclosure risk can be assessed appropriately by comparing the intruder's knowledge before and after data release; see Duncan and Lambert (1986, 1989), Lambert (1993) and Keller-McNulty et al. (2005) for measures of predictive disclosure risk, based on the predictive distribution of Y given x.

The masked data, for public release, are created by altering some of the records in the original data, to reduce disclosure risk. Disclosure control techniques dilute, suppress, and in some cases distort the information content of the original data and hence compromise the quality of statistical inferences that can be made from the released data. Random noise addition usually does not bias statistical estimates, only inflates their standard errors, making them less reliable. Grouping, cell collapsing and cell suppression hide some information and certain inferences that

can be made from the original data cannot be made from the released data. These procedures also inflate the variance of the estimates that can be obtained from the masked data. Data swapping often lead to biased estimates of correlations among the survey variables. From a logical perspective it is important to note that the utility of a masked data set is less than that of the original data set, and one should take both data utility and disclosure risk into account in selecting appropriate disclosure control techniques in specific situations. Some approaches to measuring data utility and optimizing the risk-utility tradeoff have been discussed by Duncan and Fienberg (1999), Duncan and Mukherjee (2000), Duncan and Stokes (2004) and Karr et al. (2006). In particular, Keller-McNulty et al. (2005) proposed a decision theoretic framework where the data utility and disclosure risk are formulated as the utilities of a legitimate data user and of an intruder, respectively, and both are quantified in terms of Shannon's information entropy. Then, a weighted combination of the two utilities, representing the utility of the data agency, is maximized to determine the optimal masking of the original data.

The main goal of this paper is to examine two recently introduced disclosure control techniques, viz., the perturbation method (Muralidhar and Sarathy, 2003) and the PRAM (Gouweleeuw et al., 1998). In Section 2, we describe and analyze the perturbation method. We show that for implementing the procedure one needs to estimate the conditional distribution of Y given X from the original data, in which case the past claims about disclosure avoidance and no information loss do not hold. In Section 3, we discuss PRAM and its effect on disclosure control. While PRAM is similar to randomized response surveys, we explain that respondents' privacy should be assessed differently in the two cases. For PRAM, we suggest to measure (identity) disclosure risk using the probability of correct identification of an "at risk" unit. This risk is a random variable and we contend that the choice of a specific PRAM should be made based on the risk

distribution. We hope our discussions will stimulate further refinements of these procedures to make them more sound and effective disclosure control techniques.

2. Perturbation Technique

A general perturbation procedure, described next, was proposed by Muralidhar and Sarathy (2003), who also claimed that the procedure yields masked microdata that contain the same statistical information as the original data set and do not provide any additional information to an intruder (and hence protects full confidentiality). As before, let X denote the key variables, Y denote the confidential variables and let x_i, y_i denote the values of X, Y for the ith record in the original data set. Using $f(.)$ to denote a (generic) probability density function (pdf), Muralidhar and Sarathy (2003) recommended to replace, for each record i, the value y_i by a simulated value z_i from $f(y|X = x_i)$. Then, the masked data set consist of the (x_i, z_i) values, which can be released to the public. Muralidhar and Sarathy (2003) claimed two features of the procedure: (i) $f(X, Z) = f(X, Y)$, implying that the masked data contain the same statistical information as the original data, and (ii) $f(Y|X, Z) = f(Y|X)$ and so releasing of Z does not provide any additional information about Y. In the following we examine the feasibility of the proposed procedure and the validity the preceeding claims.

Our comments are centered around the meaning of $f(.)$, which appeared in various forms in the preceeding discussion (adopted from Muralidhar and Sarathy, 2003) without adequate clarification, and we find that different interpretations of $f(.)$ lead to different conclusions. First, we shall note that the statements $f(X, Z) = f(X, Y)$ and $f(Y|X, Z) = f(Y|X)$ hold true in a particular context. For clarity, we shall denote various distributions with appropriate subscripts

of f. Suppose two vectors x and y are generated from a specific (and given) joint pdf $f^*(.,.)$ and then a vector z is generated from the conditional distribution $f^*_x(.) = f^*(x,.)/[\int f^*(x,u)du]$. Then, denoting the associated random vectors by X, Y and Z, mathematically it follows that

$$f_{X,Y}(a,b) = f_{X,Z}(a,b) = f^*(a,b) \qquad \text{for all } (a,b) \tag{2.1}$$

and

$$f_{Y|X=x,Z=z}(a) = f_{Y|X=x}(a) = f^*_x(a) \qquad \text{for all } a, x \text{ and } z. \tag{2.2}$$

Equation (2.1) shows that (X,Y) and (X,Z) contain the same statistical information about $f^*(.,.)$, and Eq. (2.2) implies that Y and Z are conditionally independent given X, i.e., when X is given, Z does not contain any additional information about Y. However, the applicability of these results to Muralidhar and Sarathy's (2003) procedure is uncertain, as we discuss in the following.

In Muralidhar and Sarathy (2003), it is not clear what $f(y|X = x_i)$ represents. So, we shall consider three possible interpretations and explore their consequences. First, suppose the data are regarded as a sample, i.e., x_i and y_i are randomly generated from a joint pdf $f_{X_i,Y_i}(.,.)$. When the data are generated by random sampling, $f_{X_i,Y_i}(.,.)$ is the same for all i, and the common distribution, to be denoted by $f_{X,Y}(.,.)$, is the population distribution, which is unknown in the context of data dissemination; if the population distribution is known, there would be no need for data. In general, $f_{X_i,Y_i}(.,.)$ may depend on i, as in the case of stratified sampling. In any case, $f_{X_i,Y_i}(.,.)$ are unknown distributions. For simplicity, consider the case of random sampling. Then, the masked variable Z would satisfy (2.1) and (2.2) (with $f^*(.,.) = f_{X,Y}(.,.)$) if for each record i, z_i is generated from $f_{Y|X=x_i}(y) = f_{X,Y}(x_i,y)/[\int f_{X,Y}(x_i,y)dy]$, i.e., in the perturbation procedure $f(y|X = x_i)$ is interpreted as the conditional pdf of Y given $X = x_i$ in

the population. This, however, is not feasible as the joint and conditional distributions for the population are unknown.

Now, suppose the original data are viewed as the records of a population, and in Muralidhar and Sarathy (2003) $f(X, Y)$ represents the distribution of X and Y in the original data set. Then, (2.1) can hold only if the masked data set $\{(x_i, z_i)\}$ contain the same set of records (possibly in a permuted order) as the original data set. To see this, suppose (for simplicity) both X and Y are categorical with possible categories $c_1, \cdots, c_k$ for X, and $d_1, \cdots, d_l$ for Y, and suppose n_{ij} is the frequency of the pair $(X = c_i, Y = d_j)$ in the original data set. Then, $f_{X,Y}(c_i, d_j) = n_{ij}/n$, where n is the total number of records in the data set. Let n^*_{ij} be the frequency of the pair $(X = c_i, Z = d_j)$ in the masked data set. Then, the joint pdf of X and Z, taking the masked data set as a population, is $f_{X,Z}(c_i, d_j) = n^*_{ij}/n$. So, Eq. (2.1) would hold if and only if $n^*_{ij} = n_{ij}$, i.e., the original and the masked data sets contain exactly the same set of records, in which case masking does nothing for confidentiality protection. In particular, (2.2) is not satisfied in this interpretation of $f(X, Y)$.

A third possible interpretation of Muralidhar and Sarathy's (2003) proposal is: treat the data (original as well as masked) as a sample, interpret $f(X, Y) = f_{X,Y}(.,.)$ as the unknown population distribution, but for $f(y|X = x_i)$, which is unknown, use an estimate of the conditional pdf of Y given X in the population, based on the original data. This is a feasible approach (and is also used in imputation techniques) but the resulting Z does not satisfy (2.1) and (2.2). So, the claims of no statistical information loss and no additional risk of disclosure do not hold for this approach. The change in statistical information, viz., the difference between $f_{X,Y}(.,.)$ and $f_{X,Z}(.,.)$, depends on the accuracy of the estimated $f(y|x)$. When X and Y are continuous variables, one is likely to estimate the conditional distributions using a parametric model, such

as multivariate normal (as nonparametric estimation may lack adequate precision). The main reason for releasing microdata is to allow researchers to explore and analyze the data using other models. However, the perturbed data generated using a fitted model may not be suitable for such purposes as the fitted model induces all features of Z and its relationship with X. So, fitting other models to the released data may yield misleading conclusions. Perhaps, data users will be better served if agencies release only the x-values and the fitted model for Y along with some summary measures of accuracy of the fitted model, or if the y-values are imputed only for a small subset of "at risk" units.

3. The Post Randomization Method

The Post Randomization Method (PRAM), introduced by Kooiman et al. (1997) and further discussed in Gouweleeuw et al. (1998) and Willenborg and de Wall (2001), is a disclosure control technique for categorical variables and it is analogous to noise addition to values of continuous variables. The procedure, when applied to a categorical variable, alters each record on that variable using a pre-selected probability mechanism. Let X be a categorical variable with categories $c_1, \cdots, c_k$. Then, PRAMing of the records on X is done by replacing the value c_l $(l = 1, \cdots, k)$ of X by c_i with probability p_{il} and independently for each record, where $p_{il}, i, l = 1, \cdots, k$ are prespecified and satisfy the conditions:

$$\sum_{l=1}^{k} p_{il} = 1, \quad i = 1, \cdots, k.$$

We shall denote the transformed variable by X^*, in which case $P(X^* = c_i | X = c_l) = p_{il}$. The PRAMed data, created by replacing all records on X by the transformed values, are released to the public along with the transition probability matrix $P = ((p_{ij}))$. The choice of P affects the

degree of disclosure control and as well as information masking.

Mathematically, PRAM is similar to randomized response (RR) surveys (see, e.g., Warner, 1965; Chaudhuri and Mukerjee, 1988; Nayak, 1994) as was noted by Gouweleeuw et al. (1998) and Van den Hout and Van der Heijden (2002). Operationally, in RR surveys each respondent randomizes his or her true response at data gathering stage, whereas in PRAM randomization is performed by the surveyor after the data are collected. Statistical methods for analyzing RR data are applicable to PRAMed data, but as we discuss later in this section, the protection of privacy needs to be assessed differently in the two contexts.

PRAM can be applied to several categorical variables, independently or jointly on the cross-classification of those variables (see, Gouweleeuw et al., 1998). Conceptually, any PRAM procedure can be regarded as applied to the combined variable created by cross-classifying all categorical variables. For any PRAM procedure, there is a known transition probability matrix for the combined variable, e.g., if the variables are PRAMed independently, the transition probability matrix for the combined variable can be expressed as a Kronecker product of the transition probability matrices for the individual variables. For estimating the joint probabilities, it is convenient (and perhaps appropriate) to consider a single variable obtained by compounding all categorical variables. Let X represent the compound variable, $\pi_i = P(X = c_i), \lambda_i = P(X^* = c_i), i = 1, \cdots, k, \Pi = (\pi_1, \cdots, \pi_k)'$ and $\Lambda = (\lambda_1, \cdots, \lambda_k)'$. Note that Π is of primary inferential interest and $\Lambda = P\Pi$. A natural estimate of λ_i is $\hat{\lambda}_i = \{\#(X^* = c_i)\}/n$, where n is the sample size. Under multinomial sampling, $\hat{\lambda}_i$ is both MLE and UMVUE of λ_i. An estimator $\hat{\Lambda}$ of Λ, based on the PRAMed data, yields the natural estimator $\hat{\Pi} = P^{-1}\hat{\Lambda}$ of Π, provided that P is nonsingular. Obviously, if $\hat{\Lambda}$ is unbiased, then so

is $\hat{\Pi}$. Under multinomial sampling it can be seen that (Chaudhuri and Mukerjee, 1988, p. 43)

$$Var(\hat{\Pi}) = \frac{(D_\Pi - \Pi\Pi')}{n} + \frac{[P^{-1}D_\Lambda(P^{-1})' - D_\Pi]}{n},$$

where D_Π is a diagonal matrix with diagonal elements being $\pi_1, \cdots, \pi_k$ and D_Λ is defined similarly. The first term on the right side is the variance under no randomization and the last term is the additional variance induced by PRAM. Since Π can be estimated unbiasedly from PRAMed data, any limitation of predictive disclosure comes from the variance inflation.

We now proceed to make some comments on the effect of PRAM on disclosure limitation using the following example from Gouweleeuw et al. (1998). Suppose a microdata file contains a simple random sample of n records from a population of size N and only the gender variable is PRAMed with $p_{ii} = .9, i = 1, 2$, i.e., for each record, the gender value is changed to the opposite gender with probability .1 and is kept unchanged with probability .9. Suppose an intruder knows that the population contains 99 male surgeons and one female surgeon. In this context, Gouweleeuw et al. (1998) state that the intruder "can derive that the probability that the female surgeon in the perturbed file is indeed the female surgeon in the population equals 0.08. This probability is very small, hence the perturbed data can be considered safe." These comments relate to identity disclosure risk and we believe, the above conclusion is derived from the fact that for a randomly selected surgeon from the population, $P(X = f|X^* = f) = 0.08$. This probability is a relevant measure of respondents' privacy in the context of a randomized response survey, where the identity of the respondent is known to the surveyor but not the value of X, and the issue is: how accurately the surveyor can predict the value of X from that of X^*. However, as we discuss below, the relevance of $P(X = f|X^* = f)$ as a measure of identity disclosure limitation is not clear.

Let S denote the number of surgeons and T denote the number of female surgeons in the masked data set. Note that both S and T are random variables, due to sampling as well as PRAMing. So, the value of T need not be exactly one. Also, $E(S) = 100(n/N)$, which depends on the sampling ratio n/N. The conditional probability of disclosure of the female surgeon's identity, given $S = s$ and $T = t$, depends on s and t. In the following we examine this probability for general s and $t = 1$. It can be seen that

$$P(S = s, T = 1) = \frac{\binom{N-100}{n-s}[\binom{99}{s-1}\{(.9)^s + (s-1)(.1)^2(.9)^{s-2}\} + \binom{99}{s}s(.1)(.9)^{s-1}]}{\binom{N}{n}}. \tag{3.1}$$

It can also be seen that the probability that $S = s, T = 1$ and the female surgeon in the masked data is the female surgeon in the population is

$$\frac{\binom{N-100}{n-s}\binom{99}{s-1}(.9)^s}{\binom{N}{n}}. \tag{3.2}$$

From (3.1) and (3.2) it follows that if the masked data show one female surgeon and $(s-1)$ male surgeons, the probability that the female surgeon is really female is $0.81/[9.8 - (0.08)s]$. This probability, which we believe appropriately measures the female surgeon's identity disclosure risk when the PRAMed data set show one female surgeon, depends on s, but interestingly, is independent of N and n. The numerical values of this probability, for s=1, 2, 10, 30 and 50, are 0.083, 0.084, 0.090, 0.109 and 0.140, respectively. In our example, the disclosure risk increases modestly with s, but a different gender distribution of the surgeon population may exhibit much greater change.

If the female surgeon is not in the sample, arguably there can be no disclosure of her information. The issue of disclosure is more meaningful when the subject is in the sample and the intruder knows that. So, let us suppose the intruder's information is about the sample. Suppose, in our example, the intruder knows the sample contains 100 surgeons one of whom is female. In

Table 1: Probability of correct match

t	$P(t)$	$P(\text{correct match}\|t)$	t	$P(t)$	$P(\text{correct match}\|t)$
1	.00006	.4500	13	.0941	.0711
2	.0005	.3115	14	.0695	.0664
3	.0022	.2382	15	.0472	.0623
4	.0074	.1929	16	.0296	.0587
5	.0188	.1620	17	.0172	.0555
6	.0384	.1397	18	.0093	.0526
7	.0652	.1227	19	.0047	.0500
8	.0944	.1095	20	.0022	.0476
9	.1188	.0988	21	.0010	.0455
10	.1319	.0900	22	.0004	.0435
11	.1305	.0827	23	.00016	.0418
12	.1164	.0764	24	.00006	.0401

this case S is fixed (at 100) but T is random (due to PRAMing) and $E(T) = 1(.9)+99(.1) = 10.8$. For identification of the female surgeon's record, suppose the intruder randomly selects one of the T records in masked data set that are shown as female surgeons. Then, the probability that the intruder gets a correct match, given that $T = t > 0$, can be seen to be

$$P(\text{correct match}|T = t) = \frac{0.81}{1 + (0.8)t}. \tag{3.3}$$

In Table 1, we present the probability distribution of T and the conditional probability of correct match given $T = t$. While the possible values of t are $0, 1, \cdots, 100$, in Table 1, we present only $t = 1, \cdots, 24$ as the other values have very small probability. As t increases from 1 to 24, the conditional probability of correct match decreases from .45 to .04. The most likely value of T is

10, in which case the probability of correct match is 0.09. The value of t, which determines the disclosure risk, is generated after PRAM is applied. On the other hand, at the time of selecting the transition probability matrix P, the value of t is not known. So, for evaluating and selecting a P, one should look at both the probability distribution of T and the conditional probabilities of correct match. A conservative approach would be to consider all values of t with probability exceeding a given threshold (α), as the "likely" values of t, and then look at the maximum match probability over that set. As an illustration, in our example, the maximum match probability for $\alpha = 0.02$ is 0.1397 (corresponding to $t = 6$), which is not too small. This suggests that for adequate disclosure control we should change the gender value with a probability higher than 0.9, the current value. For adequate disclosure control for the subjects in a sensitive category of X, say $X = c$, Gouweleeuw et al. (1998) suggested to keep the posterior odds

$$PO(c) = \frac{P(X = c|X^* = c)}{P(X \neq c|X^* = c)}$$

lower than a selected value, say β. This is equivalent to keeping $P(X = c|X^* = c)$ smaller than $\beta/(1 + \beta)$. This approach is simple and convenient (and appropriate for ensuring respondents' privacy in randomized response surveys), but may not be satisfactory for disclosure limitation, as discussed above. The choice of the transition probabilities, for controlling disclosure risk, should more appropriately be made based on the entire risk distribution. We hope the conservative approach suggested above will be useful in many practical situations.

Acknowledgment. This research was supported in part by a grant from the Institute for Integrating Statistics in Decision Sciences at The George Washington University.

References

[1] Bethlehem, J.G., Keller, W.J. and Pannekoek, J. (1990). Disclosure control of microdata. *J. Amer. Statist. Assoc.*, 85, 38-45.

[2] Chaudhuri, A. and Mukerjee, R. (1988). *Randomized Response: Theory and Techniques.* Marcel Dekker, New York.

[3] Doyle, P., Lane, J., Theeuwes, J. and Zayatz, L. (Ed.) (2001). *Confidentiality, Disclosure and Data Access: Theory and Practical Applications for Statistical Agencies.* Amsterdam: Elsevier.

[4] Duncan, G.T. and Fienberg, S.E. (1999). Obtaining information while preserving privacy: A Markov perturbation method for tabular data. in *Eurostat Statistical Data Protection '98 Lisbon*, Luxembourg: Eurostat, pp. 351-362.

[5] Duncan, G.T. and Lambert, D. (1986). Disclosure-limited data dissemination. *J. Amer. Statist. Assoc.*, 81(393), 10-19.

[6] Duncan, G.T. and Lambert, D. (1989). The risk of disclosure for microdata. *J. Business & Econ. Statist*, 7, 207-217.

[7] Duncan, G.T. and Stokes, S.L. (2004). Disclosure risk vs. data utility: The R-U confidentiality map as applied to topcoding. *Chance*, 17, 16-20.

[8] Duncan, G. T. and Mukherjee, S. (2000) Optimal disclosure limitation strategy in statistical databases: Deterring tracker attacks through additive noise. *J. Amer. Statist. Assoc.*, 95(451):720–729.

[9] Gouweleeuw, J.M., Kooiman, P., Willenborg, L.C.R.J. and De Wolf, P.P. (1998). Post randomisation for statistical disclosure control: Theory and implementation. *J. Official Statist.*, 14, 463–478.

[10] Greenberg, B. and Zayatz, L. (1992). Strategies for measuring risk in public use microdata files. *Statist. Neerland.*, 46, 33-48.

[11] Keller-McNulty, S., Nakhleh, C.W. and Singpurwalla, N.D. (2005). A Paradigm for Masking (Camouflaging) Information. *Internat. Statist. Rev.*, 73, 331-349.

[12] Kooiman, P., Willenborg, L., and Gouweleeuw, J. (1997). A method for disclosure limitation of microdata. Research paper 9705, Statistics Netherlands, Voorburg.

[13] Lambert, D. (1993). Measure of disclosure risk and harm. *Journal of Official Statistics*, 9(2), 313331.

[14] Muralidhar, K. and Sarathy, R. (2003). A theoretical basis for perturbation methods. *Statistics and Computing*, 13(4), 329-335.

[15] Muralidhar, K. and Sarathy, R. (2006). Data shuffling - A new masking approach for numerical data. *Management Science*, 52(5), 658-670.

[16] Nayak, T.K. (1994). On Randomized Response Surveys for Estimating a Proportion. *Commun. Statist.- Theory Meth.*, 23(11), 3303-3321.

[17] Reiter, J.P. (2005). Estimating identification risk in microdata. *J. Amer. Statist. Assoc.*, 100, 1101-1113.

[18] Skinner, C.J. and Elliot, M.J. (2002). A measure of disclosure risk for microdata. *J. R. Statist. Soc.*, Ser. B, 64, 855-867.

[19] Van den Hout, A., and van der Heijden, P.G.M. (2002). Randomized response, statistical disclosure control and misclassification: a review. *Internat. Statist. Rev.*, 70, 269-288.

[20] Warner, S.L. (1965). Randomized response: a survey technique for eliminating evasive answer bias. *J. Amer. Statist Assoc.*, 60, 63-69.

[21] Willenborg, L.C.R.J. and De Waal, T. (2001). *Elements of Statistical Disclosure Control*. New York: Springer.