Statistical Disclosure Control using Post Randomisation: Variants and Measures for Disclosure Risk

Ardo van den Hout, Elsayed A. H. Elamir

Abstract

This paper discusses the post randomisation method (PRAM) as a method for disclosure control. PRAM protects the privacy of respondent by misclassifying specific variables before data are released to researchers outside the statistical agency. Two variants of the initial idea of PRAM are discussed concerning the information about the misclassification that is given along with the released data. The first variant concerns calibration probabilities and the second variant concerns misclassification proportions. The paper shows that the distinction between the univariate case and multivariate case is important. Additionally, the paper discusses two measures for disclosure risk when PRAM is applied.

S³RI Methodology Working Paper M04/01

Statistical Disclosure Control Using Post Randomisation: Variants and Measures for Disclosure Risk

Ardo van den Hout

Department of Methodology and Statistics
Faculty of Social Sciences, Utrecht University
Heidelberglaan 1, 3584 CS Utrecht, the Netherlands
tel. +31 30 2539237, E-mail: A.vandenHout@fss.uu.nl

Elsayed A.H. Elamir

Southampton Statistical Sciences Research Institute
University of Southampton
Southampton SO17 1BJ, United Kingdom
tel: +44 23 8059 4298, E-mail: E.H.Elamir@socsci.soton.ac.uk

Abstract

This paper discusses the post randomisation method (PRAM) as a method for disclosure control. PRAM protects the privacy of respondent by misclassifying specific variables before data are released to researchers outside the statistical agency. Two variants of the initial idea of PRAM are discussed concerning the information about the misclassification that is given along with the released data. The first variant concerns calibration probabilities and the second variant concerns misclassification proportions. The paper shows that the distinction between the univariate case and the multivariate case is important. Additionally, the paper discusses two measures for disclosure risk when PRAM is applied.

Keywords: calibration; information loss; misclassification; PRAM.

1 Introduction

The post randomisation method (PRAM) is discussed in Gouweleeuw, Kooiman, Willenborg, and De Wolf (1998) as a method for statistical disclosure control (SDC). When survey data are released by statistical agencies, SDC protects the identity of respondents. SDC tries to prevent that a user of the released data can link the data of a respondent in the survey to a specific person in the population. See Willenborg and De Waal (2001) for an introduction into SDC and SDC methods other than PRAM.

There is a close link between PRAM and randomised response, a method to ask sensitive questions in a survey, see Warner (1965) and Rosenberg (1979). Van den Hout and Van der Heijden (2002) sum up some differences and similarities between randomised response and PRAM.

When SDC is used, there will always be a loss of information. This is inevitable since SDC tries to determine the information in the data that can lead to the disclosure of an identity of a respondent, and eliminates this information before data are released. It is not difficult to prevent disclosure, but it is difficult to prevent disclosure *and* release data that is still useful for statistical analysis. Applying SDC means searching for a balance between disclosure risk and information loss.

The idea of PRAM is to misclassify some of the categorical variables in the survey using fixed misclassification probabilities and to release the partly misclassified data together with those probabilities. Say variable X, with categories $\{1, ..., J\}$, is misclassified into variable X^*. The survey containing X^* but not X is released together with conditional probabilities $\mathbb{P}(X^* = k|X = j)$ for $k, j \in \{1, ..., J\}$. In this way PRAM introduces uncertainty in the data: the user of the released data cannot be sure that the information is original or perturbed due to PRAM and it becomes harder to establish a correct link between a respondent in the survey and a specific person in the population. Since the user has the misclassification probabilities he can adjust his analysis by taking into account the perturbation due to PRAM.

This paper discusses two ideas to make PRAM more efficient with respect to the balance between disclosure risk and information loss. First, the paper discusses the use of *calibration probabilities*

$$\mathbb{P}(\text{true category is } j|\text{category } i \text{ is released}). \tag{1}$$

in the analysis of released data and compares this with using *misclassification probabilities*

$$\mathbb{P}(\text{category } i \text{ is released}|\text{true category is } j). \tag{2}$$

The idea of using calibration probabilities is discussed by De Wolf, Gouweleeuw, Kooiman, and Willenborg (1997), who refer to the discussion of calibration probabilities in misclassification literature, see, e.g., Kuha and Skinner (1997). We will elaborate the discussion and show that the advantage of calibration probabilities is limited to the univariate case. Secondly, the paper shows that information loss can be reduced by providing *misclassification proportions* along with the released data. These proportions inform about the actual change in the survey data due to the application of PRAM. (Probabilities (1) and (2) inform about the expected change.) Additionally, the paper discusses two measures for disclosure risk when PRAM is applied. The first is a general measure presented in Elamir and Skinner (2003) as an extension of the measure introduced by Skinner and Elliot (2002). The second measure links up with the SDC practice at Statistic Netherlands. Simulation results are given to illustrate the theory.

The outline of the paper is as follows. Section 2 provides the framework and the notation. Section 3 describes frequency estimation for PRAM data. Section 4 discusses the use of calibration probabilities. In Section 5 we introduce the use of misclassification proportions. Section 6 discusses measures for disclosure risk, whereas information loss is briefly considered in Section 7. Section 8 presents some simulations, and Section 9 concludes.

2 Framework and notation

In survey data we distinguish between *identifying variables* and *non-identifying variables.* Identifying variables are variables that can be used to re-identify individuals represented in the data. These variables are assumed to be categorical, e.g., Gender, Race, Place of Residence. We assume that the sensitive information of respondents is contained in the non-identifying variables, see Bethlehem, Keller, and Pannekoek (1990), and that we want to protect this information by applying PRAM to (a subset of) the identifying variables.

The notation in this paper is the same as in Skinner and Elliot (2002). Units are selected from a finite population U and each selected unit has one record in the *microdata sample* $s \subset U$. Let n denote the number of units in s. Let the categorical variable formed by cross-classifying (a subset of) the identifying variables be denoted X with values in $\{1, ..., J\}$. Let X_i denote the value of X for unit $i \in U$. The *population frequencies* are denoted

$$F_j = \sum_{i \in U} I(X_i = j), \qquad j \in \{1, ..., J\},$$

where $I(\cdot)$ is the indicator function: $I(A) = 1$ if A is true and $I(A) = 0$ otherwise.

The *sample frequencies* are denoted

$$f_j = \sum_{i \in s} I(X_i = j), \qquad j \in \{1, ..., J\}.$$

In the framework of PRAM, we call the sample that is released by the statistical agency the *released microdata sample* s^*. Note that unit $i \in s^*$ if and only if $i \in s$. Let X^* denote the released version of X in s^*. By *misclassification* of unit i we mean $X_i \neq X_i^*$. The *released sample frequencies* are denoted

$$f_k^* = \sum_{i \in s^*} I(X_i^* = k), \qquad k \in \{1, ..., J\}.$$

Let P_X denote the $J \times J$ transition matrix that contains the conditional misclassification probabilities $p_{kj} = I\!P(X^* = k|X = j)$, for $k, j \in \{1, ..., J\}$. Note that the columns of P_X sum up to one. The distribution of X^* conditional on s is the J-component finite mixture given by

$$I\!P(X_i^* = k|i \in s) = \sum_{j=1}^{J} I\!P(X_i^* = k|X_i = j) I\!P(X_i = j|i \in s), \qquad k \in \{1, ..., J\},$$

where the component distributions are given by P_X and the component weights are given by the conditional distribution of X. The conditional distribution of X in sample s is multinomial with

$$I\!P(X_i = j|i \in s) = \frac{1}{n} f_j, \qquad j \in \{1, ..., J\}.$$

3 Frequency estimation for PRAM data

When PRAM is applied and some of the identifying variables are misclassified, standard statistical models do not apply to the released data since these models do not take into account the perturbation. This section shows how the misclassification can be taken into account in frequency estimation.

We have $\mathbb{E}[\boldsymbol{F}^*|\boldsymbol{f}] = P_X\boldsymbol{f}$, where $\boldsymbol{f} = (f_1, ..., f_J)^t$ and $\boldsymbol{F}^* = (F_1^*, ..., F_J^*)^t$ is the stochastic vector of the released sample frequencies. An unbiased moment estimator of $\boldsymbol{f}$ is given by

$$\hat{\boldsymbol{f}} = P_X^{-1}\boldsymbol{f}^*, \tag{3}$$

see Kooiman, Willenborg, and Gouweleeuw (1997). In practice, assuming that P_X is non-singular does not impose much restriction on the choice of the misclassification probabilities. Matrix P_X^{-1} exists when the diagonal of P_X dominates, i.e., $p_{ii} > 1/2$ for $i \in \{1, ..., J\}$. An additional assumption is that the dimensions of $\boldsymbol{f}$ and $\boldsymbol{f}^*$ are the same.

PRAM is applied to each variable independently and a transition matrix is released per variable. When the user of the released sample assesses a compounded variable, he can construct its transition matrix using the transition matrices of the individual variables. For instance, consider identifying variables X_1, with categories $\{1, .., J_1\}$ and X_2, with categories $\{1, .., J_2\}$, and the cross-classification $X = (X_1, X_2)$, i.e., the Cartesian product of X_1 and X_2. Since PRAM is applied independently, we have

$$\begin{aligned} \mathbb{P}\Big(X^* = (k_1, k_2)|X = (j_1, j_2)\Big) &= \mathbb{P}(X_1^* = k_1|X_1 = j_1) \\ &\times \ \mathbb{P}(X_2^* = k_2|X_1 = j_2), \end{aligned} \tag{4}$$

for $k_1, j_1 \in \{1, .., J_1\}$ and $k_2, j_2 \in \{1, .., J_2\}$. In matrix notation, we have $P_X = P_{X_1} \otimes P_{X_2}$, where $\otimes$ is the Kronecker product. Note that when one of two variables is not perturbed by PRAM, the transition matrix of that variable is the identity matrix.

The variance of (3) equals

$$V[\hat{\boldsymbol{f}}|\boldsymbol{f}] = P_X^{-1}V[\boldsymbol{F}^*|\boldsymbol{f}](P_X^{-1})^t = P_X^{-1}\Big(\sum_{j=1}^{J} f_j V_j\Big)(P_X^{-1})^t \tag{5}$$

where V_j is the $J \times J$ covariance matrix of two released values given the original value j, i.e.,

$$V_j(k_1, k_2) = \begin{cases} p_{k_2 j}(1 - p_{k_2 j}) & \text{if} \quad k_1 = k_2 \\ \\ -p_{k_1 j} p_{k_2 j} & \text{if} \quad k_1 \neq k_2 \end{cases} \qquad \text{for } k_1, k_2 \in \{1, ..., J\},$$

see Kooiman *et al.* (1997). The variance can be estimated by substituting $\hat{f}_j$ for f_j in (5), for $j \in \{1, .., J\}$.

The variance given by (5) is the extra variance due to PRAM and does not take into account the sampling distribution. The formulas for the latter are given in Chaudhuri and Mukerjee (1988) for multinomial sampling and compared to (5) in Van den Hout and Van der Heijden (2002), see also Appendix B.

4 Calibration probabilities

Literature concerning misclassification shows that calibration probabilities (1) are more efficient in the analysis of misclassified data than misclassification probabilities (2), see the review paper by Kuha and Skinner (1997). Often, calibration probabilities have to be estimated. However, when PRAM is applied, the statistical agency can compute the calibration probabilities using the sample frequencies. The idea of using calibration probabilities for PRAM is mentioned in De Wolf *et al.* (1997). The following elaborates this idea and makes a comparison with PRAM as explained in the previous section.

The $J \times J$ matrix with calibration probabilities of univariate variable X is denoted by $\overleftarrow{P}_X$ and has entries $\overleftarrow{p}_{jk}$ defined by

$$\mathbb{P}(X_i = j | X_i^* = k, i \in s) = \frac{p_{kj} f_j}{\sum_{j_0=1}^{J} p_{kj_0} f_{j_0}}, \qquad j, k \in \{1, ..., J\}, \qquad (6)$$

where p_{kj} are the entries of P_X. Matrix $\overleftarrow{P}_X$ is again a transition matrix; each column sums up to one. We have

$$\boldsymbol{f} = \overleftarrow{P_X}\mathbb{E}[\boldsymbol{F}^*|\boldsymbol{f}], \tag{7}$$

see Appendix A. An unbiased moment estimator of $\boldsymbol{f}$ is therefore given by

$$\tilde{\boldsymbol{f}} = \overleftarrow{P}_X \boldsymbol{f}^*. \tag{8}$$

In general, $\overleftarrow{P}_X \neq P_X^{-1}$, see Appendix A. The variance of (8) is given by (5) where P_X^{-1} is replaced by $\overleftarrow{P}_X$ and f_j is estimated by $\tilde{f}_j$, for $j \in \{1, ..., J\}$.

In the remainder of this section we compare estimators (3) and (8). The first difference is that (3) might yield an estimate where some of the entries are negative, whereas (8) will never yield negative estimates, see, e.g., De Wolf *et al.* (1997).

Secondly, estimator (8) is more efficient than (3) in the univariate case. This is already discussed in Kuha and Skinner (1997). Consider the case where X has two categories. Say we want to know $\pi = \mathbb{P}(X = 1)$. Let $\hat{\pi}$ be the estimate using P_X and $\tilde{\pi}$ the estimate using $\overleftarrow{P}_X$. The efficiency of $\hat{\pi}$ relative to $\tilde{\pi}$ is given by

$$\text{eff}(\hat{\pi}, \tilde{\pi}) = \frac{V[\tilde{p}\,]}{V[\hat{p}\,]} = (p_{11} + p_{22} - 1)^2(\overleftarrow{p}_{22} - \overleftarrow{p}_{21})^2 < 1. \tag{9}$$

So $\tilde{\pi}$ is always more efficient than $\hat{\pi}$. An important difference with the general situation of misclassification is that in the situation of PRAM, matrices P_X and $\overleftarrow{P}_X$ are given and do not have to be estimated. Comparison (9) is therefore a simple form of the comparison in Kuha and Skinner (1997, Section 28.5.1.3.).

The third comparison is between the maximum likelihood properties of (3) and (8). Assume that $X_1, ...X_n$ are independently multinomially distributed with parameter vector $\boldsymbol{\pi} = (\pi_1, .., \pi_J)^t$. In the framework of misclassification, Hochberg (1977) proves that estimator (8) yields a MLE. When (3) yields an estimate in the interior of the parameter space, the estimate is also a MLE. See Appendix B for the maximum likelihood properties of (8) and (3). Note that the likelihood function corresponding to (8) is different from the likelihood function corresponding to (3), since the information used is different. This explains why both can be MLE despite being different estimators of $\boldsymbol{f}$.

The fourth comparison is with respect to transition matrices of Cartesian products and is less favourable for (8). It has already been noted that $P_{X_1} \otimes P_{X_2}$ is the matrix with misclassification probabilities for the Cartesian product $X = (X_1, X_2)$, see (4). Analogously, given $\overleftarrow{P}_{X_1}$ and $\overleftarrow{P}_{X_2}$ the user can construct matrix $\overleftarrow{P}_{X_1} \otimes \overleftarrow{P}_{X_2}$. However, this matrix does *not* necessarily contain calibration probabilities for X. Note that

$$\mathbb{P}\Big(X_i = (j_1, j_2) | X_i^* = (k_1, k_2), i \in s\Big) \\ = \frac{p_{k_1 j_1} p_{k_2 j_2} \mathbb{P}\Big(X_i = (j_1, j_2) | i \in s\Big)}{\sum_v^{J_1} \sum_w^{J_2} p_{k_1 v} p_{k_2 w} \mathbb{P}\Big(X_1 = (v, w) | i \in s\Big)}, \tag{10}$$

It follows that $\overleftarrow{P}_X = \overleftarrow{P}_{X_1} \otimes \overleftarrow{P}_{X_2}$ when X_1 and X_2 are independent. In general, this independence is not guaranteed and since the user of the released data does not have the frequencies of X, he cannot construct $\overleftarrow{P}_X$.

The fifth and last comparison is with respect to the creation of subgroups. Consider the situation where a user of the released data creates a subgroup by using a grouping variable that is not part of X. When the number of categories in the subgroup is smaller than J, estimate (8) is not well-defined. When the number of categories is equal to J, estimate (8) is biased due to the fact that (7) does not hold. Note with respect to (7) that the frequencies that are used to construct $\overleftarrow{P}_X$ are the frequencies in the whole sample which will differ from the frequencies in the subgroup, see also Appendix A. The estimator (3) is still valid for the subgroup.

Since calibration probabilities contain information about the distribution of the sample s, they perform better than misclassification probabilities regarding the univariate case. However, in a multivariate setting this advantage may disappear.

Table 1: *Classification by X^* and X*

	X		
X^*	1	2	total
1	300	200	500
2	100	400	500
total	400	600	1000

Section 8 presents some simulation results.

5 Misclassification proportions

Matrices P_X and $\overleftarrow{P}_X$ inform about the expected change due to PRAM. As an alternative we can create transition matrices that inform about the actual change due an application of PRAM. These matrices contain proportions and will be denoted P_X° and $\overleftarrow{P}_X^\circ$. Matrix P_X° contains *misclassification proportions* and $\overleftarrow{P}_X^\circ$ contains *calibration proportions*. This section shows how P_X° and $\overleftarrow{P}_X^\circ$ are computed and discusses properties of these matrices.

We start with an example. Say that X has categories {1,2}. Assume that applying PRAM yields the cross-classification in Table 1. From this table it follows that the proportion of records with $X = 1$ that have $X^* = 1$ in the released sample is 300/400=3/4 and that the proportion of records with $X^* = 1$ that have $X = 1$ in the original sample is 300/500=3/5. Analogously we get the other entries of

$$P_X^\circ = \begin{pmatrix} 3/4 & 1/3 \\ 1/4 & 2/3 \end{pmatrix} \qquad \text{and} \qquad \overleftarrow{P}_X^\circ = \begin{pmatrix} 3/5 & 1/5 \\ 2/5 & 4/5 \end{pmatrix}.$$

For the general construction of P_X° and $\overleftarrow{P}_X^\circ$, let the cell frequencies in the cross-classification X^* by X be denoted c_{kj}, for $k, j \in \{1, .., J\}$. The entries of the $J \times J$ transition matrices with the proportions are given by

$$p_{kj}^e = \frac{c_{kj}}{f_j} \qquad \text{and} \qquad \overleftarrow{p}_{jk}^e = \frac{c_{kj}}{f_k^*},$$

where $k, j \in \{1, .., J\}$.

It follows that $\boldsymbol{f}^* = P^{\circ}_{X}\boldsymbol{f}$ and $\boldsymbol{f} = \overleftarrow{P}^{\circ}_{X}\boldsymbol{f}^*$. This is the reason to consider the matrices with the proportions more closely, since it is a great improvement compared to (3) and (8). Note that when the user of the released sample has P°_{X} or $\overleftarrow{P}^{\circ}_{X}$, he can reconstruct Table 1.

Conditional on $\boldsymbol{f}$, P°_{X} and $\overleftarrow{P}^{\circ}_{X}$ are stochastic, whereas P_X and $\overleftarrow{P}_X$ are not. In expectation P°_{X} equals P_X, and $P^{\circ}_{X_1} \otimes P^{\circ}_{X}$ equals $P_{X_1} \otimes P_{X_2}$, see appendix C. However, since $\boldsymbol{f}^*$ is a value of the stochastic vector $\boldsymbol{F}^*$, and $I\!P(F^*_k = 0) \neq 0$, the expectation of $\overleftarrow{P}^{\circ}_{X}$ does not exists. Nevertheless, an approximation shows that $\overleftarrow{P}^{\circ}_{X}$ will be close to $\overleftarrow{P}_X$, see appendix C.

There is a set back with respect to the use of proportions for Cartesian products and this is comparable to the problem mentioned in the previous section. Given $P^{\circ}_{X_1}$ and $P^{\circ}_{X_2}$ the user can construct $P^{\circ}_{X_1} \otimes P^{\circ}_{X_2}$ for $X = (X_1, X_2)$. However, $P^{\circ}_{X_1} \otimes P^{\circ}_{X_2}$ does *not* contain proportions as defined above. Note that the user does not have the cross-classification of X and X^*, so he cannot derive the proportions in P°_{X}. The same holds for $\overleftarrow{P}^{\circ}_{X}$. The optimal use of misclassification proportions is thereby limited to the univariate case.

Since misclassification proportions contain information about the actual perturbation due to PRAM, we expect them to perform well also in the multivariate case. Section 8 discusses a multivariate example.

6 Disclosure risk

There are several ways to measure disclosure risk, see, e.g., Skinner and Elliot (2002), and Domingo-Ferrer and Torra (2001). This section discusses two measures for disclosure risk with respect to PRAM. Section 6.1 discusses an extension of the general measure of disclosure risk introduced by Skinner and Elliot (2002). Section 6.2 introduces a measure that links up with the way disclosure risk is assessed at

Statistics Netherlands.

6.1 The measure θ

The following describes how the general measure for disclosure risk introduced in Skinner and Elliot (2002) can be extended to the situation where PRAM is applied before data are released by the statistical agency. When a disclosure control method such as PRAM has been applied, a measure for disclosure risk is needed to quantify the protection that is offered by the control method. Scenarios that may lead to a disclosure of the identity of a respondent of are about persons that aim at disclosure and that may have data that overlap the released data. A common scenario is that a person has a sample from another source and tries to identify respondents in the released sample by matching records. Using the extension of the measure in Skinner and Elliot (2002) we can investigate how applying PRAM reduces the disclosure risk.

Under simple random sampling Skinner and Elliot (2002) introduced the measure of disclosure risk $\theta = \mathbb{P}(\text{correct match}|\text{unique match})$ as

$$\theta = \sum_j I(f_j = 1) \Big/ \sum_j F_j \; I(f_j = 1),$$

where the summations are over $j = 1, ..., J$. The measure θ is the proportion of correct matches among those population units which match a sample unique. The measure is sample dependent and a distribution-free prediction is given by

$$\widehat{\theta} = \pi n_1 \Big/ \Big(\pi n_1 + 2(1 - \pi) n_2\Big),$$

where π is the sampling fraction, $n_1 = \sum_j I(f_j = 1)$ is the number of uniques and $n_2 = \sum_j I(f_j = 2)$ is the number of twins in the sample, see Skinner and Elliot (2002). Elamir and Skinner (2003) extended θ for the situation where misclassification occurs. The extension is given by

$$\theta_{mm} = \sum_{i \in s} I(f_{X_i} = 1, X_i^* = X_i) \Big/ \sum_j F_j \; I(f_j = 1)$$

and its distribution-free prediction is given by

$$\widehat{\theta}_{mm} = \pi \sum_j \mathrm{I}(f_j = 1) P_{jj} \Big/ \Big(\pi n_1 + 2\,(1 - \pi)\, n_2 \Big),$$

where P_{jj} is the diagonal entry (j, j) of the transition matrix P which describes the misclassification, see Elamir and Skinner(2003)

Section 8 presents some simulation results with respect to the measure θ before applying PRAM, and θ_{mm} after applying PRAM.

6.2 Spontaneous recognition

Statistics Netherlands releases data in several ways. One way is the releasing of detailed survey data under contract, i.e., data are released to bona fide research institutes that sign an agreement in which they promise not to look for disclosure explicitly, e.g, by matching the data to other data files. In this situation, SDC only concerns the protection against what is called *spontaneous recognition.* This section introduces a measure for disclosure risk for PRAM data that is specific to the control for spontaneous recognition.

Controlling for spontaneous recognition means that one should prevent that certain records attract attention. A record may attract attention when a low dimensional combination of its values has a low frequency. Also, without cross-classifying, a record may attract attention when one of its values is recognized as being very rare in the population. Combinations of values with low frequencies in the sample are called *unsafe combinations.*

Statistics Netherlands uses the rule of thumb that a recognition of a combination of values of more than three variables is not spontaneous anymore. For this reason, only combinations of three variables are assessed with respect to disclosure control for spontaneous recognition.

Note that applying PRAM causes two kinds of modifications in the sample that make disclosure more difficult. First, it is possible that unsafe combinations in the sample change into apparently safe ones in the released sample, and, secondly, it is possible that safe combinations in the sample change into apparently unsafe ones. Since misclassification probabilities are not that large (in order to keep analysis of the released sample possible) and the frequency of unsafe combinations is typically low, the effect of the first modification is negligible in expectation. The second modification is more likely to protect an unsafe combination j when there are a lot of combinations k, $k \neq j$, which are misclassified into j. This is the reason to focus, for a given record i with the unsafe combination of scores j, on the calibration probability

$$\mu = I\!P(X_i = j | X_i^* = j, i \in s).$$

When there are hardly any k, $k \neq j$, misclassified into j, this probability will be large, and, as a consequence, the record is unsafe. Note that combinations with frequency equal to zero are never unsafe.

Measure μ is a simplification since it ignores possible correlation between X and other variables in the sample. Note that X will be a Cartesian product and that the statistical agency can compute the calibration probability μ using (6) since the agency has the frequencies of X.

7 Information loss

Since we stressed in the introduction that SDC means searching for a balance between disclosure risk and information loss, this section indicates ways to investigate information loss due to PRAM.

First, the transition matrix P_X gives an idea of the loss of information. The more this matrix resembles the identity matrix, the less information gets lost. In general, this requires a definition of a distance between two matrices. However, we can apply PRAM using matrices that are parameterized by one parameter, denoted p_d. The idea is as follows. Each time PRAM is applied, the diagonal probabilities in the transition matrices are fixed and equal to p_d for all selected variables. In the columns, the probability mass $1 - p_d$ is equally divided over the entries that are not diagonal entries. In this situation $1 - p_d$ is a measure for the deviation from the identity matrix.

Although transition matrices give an idea of the information loss, it is hard to have an intuition about how a certain deviation from the identity matrix affects analysis of the released data. A second way to investigate information loss is the comparison of extra variances due to PRAM with respect to frequency estimation. The idea here is that when this extra variance is already substantial, more complex analyses of the released sample is probably not possible. The variance with respect to frequency estimation can be estimated using (5).

The next section will assess information loss due to PRAM using these two approaches.

8 Simulation results

The objective of this section is to illustrate the theory in the foregoing sections and to investigate disclosure risk and information loss for different choices of misclassification parameters. The population is chosen to consist of units with complete records in the British Household Survey 1996-1999. We have $N = 16710$ and we distinguish 5 identifying variables with respect to the household owner: Sex (S), Marital Status (M), Economic status (D), Socio-Economic Group (E), and Age (A), with number of categories 2, 7, 10, 7, and 8 respectively. In the following we consider simple random sampling without replacement from the population where the sample fraction π is equal to 0.05, 0.10 or 0.15. The three samples are denoted s_1, s_2 and s_3 and have sample sizes 836, 1671 and 2506 respectively.

The transition matrices used to apply PRAM to the selected variables are mostly of a simple form and determined by one parameter p_d, as described in Section 7. A more sophisticated construction of the transition matrices can reduce the disclosure risk further. An example of this fine-tuning will be given.

8.1 Disclosure risk and the measure θ

The following discusses disclosure risk by comparing the measure θ before PRAM is applied with the measure θ_{mm} after PRAM has been applied, see Section 6.1. The identifying variables are described by $X = (S, M, D, E, A)$ with $J = 7840$ possible categories.

Since the population is known, we can compute the measures and using the samples we can compute their predictions. Table 2 presents the simulation results using simple random sampling without replacement using different sampling fractions π and different choices of p_d. Given a choice of π and p_d, drawing the sample and applying PRAM is 100 times simulated. The means of the computed and predicted measures are reported in Table 2. Note that θ and $\widehat{\theta}$ reflect the risk before applying PRAM and θ_{mm} and $\widehat{\theta}_{mm}$ reflect the risk after applying PRAM.

It is clear from Table 2 that applying PRAM will reduce the risk, for example, when $p_d = 0.80$ and $\pi = 0.10$, applying PRAM reduces the risk from $\theta = 0.166$ to $\theta_{mm} = 0.055$. When p_d decreases, disclosure risk decreases too, as one might expect. Note that disclosure risk increases when sample size increases. In a larger sample, a record with a unique combination of scores is more likely to be a population unique and therefore the danger of a correct match is higher.

Table 2: *Simulation results of disclosure risk measures for* $X = (S, M, D, E, A)$ *before and after applying PRAM with* p_d.

p_d	π	θ	$\widehat{\theta}$	θ_{mm}	$\widehat{\theta}_{mm}$
	0.05	0.084	0.087	0.061	0.065
0.95	0.10	0.151	0.147	0.112	0.110
	0.15	0.217	0.215	0.165	0.166
	0.05	0.087	0.086	0.047	0.051
0.90	0.10	0.148	0.157	0.083	0.087
	0.15	0.213	0.216	0.123	0.127
	0.05	0.087	0.090	0.028	0.031
0.80	0.10	0.166	0.151	0.055	0.054
	0.15	0.213	0.221	0.065	0.064

8.2 Disclosure risk and spontaneous recognition

This following illustrates the measure μ for disclosure risk for spontaneous recognition that is discussed in Section 6.2.

Spontaneous recognition is defined for combinations up to three identifying variables, see Section 6.1. So there are 10 groups to consider. We will discuss only one of them, namely the group defined by $X = (M, D, E)$. The number of categories of X is 490. The measure for disclosure risk is given by

$$\mu = I\!P\Big((M, D, E) = (m, d, e)\Big|(M^*, D^*, E^*) = (m, d, e)\Big),$$

for those combinations of values (m, d, e) that have frequency 1 in sample s_1, s_2 or s_3. Note that when PRAM is not applied, $\mu = 1$. Table 3 shows results with respect to the maximum of μ when PRAM is applied to M, D and E. With respect to $X = (M, D, E)$ the number of unique combinations in s_1, s_2 or s_3 are 48, 44, and 53 respectively.

We draw two conclusions from the results. First, the results illustrate that the probability p_d matters, as one might expect. Second, the results show that the

Table 3: *Maximum of μ for values of (M, D, E) with frequency 1 when applying PRAM to M, D and E.*

	p_d					
	0.95	0.90	0.85	0.80	0.70	0.60
sample						
s_1 with n= 836	0.94	0.86	0.76	0.65	0.43	0.24
s_2 with n= 1671	0.94	0.85	0.74	0.61	0.36	0.18
s_3 with n= 2506	0.93	0.83	0.69	0.55	0.31	0.15

size of the sample is important. In order to protect an unsafe combination j, it is necessary that there are a lot of combinations that can change into j due to PRAM. Note that this is the other way around compared to the measure θ where a larger sample size causes a higher disclosure risk. This difference shows that different concepts of disclosure induce different methods for disclosure control.

The following introduces a method to fine-tune a transition matrix and shows that this can help to diminish the disclosure risk. The idea is to adjust one or more columns in the transition matrix of each variable that is part of an unsafe combination. Consider P_{X_1} where variable X_1 has J_1 categories. The column that is chosen first corresponds to the category of X_1 with the highest frequency in sample s, say column j. Let furthermore k be the number that corresponds to the category of X_1 with the lowest frequency in s. The columns of P_{X_1} that are not column j are constructed as explained in Section 7: p_d on the diagonal and $(1-p_d)$ equally divided over the other entries. Column j is fine-tuned by

$$p_{lj} = \begin{cases} p_d & \text{if} \quad l = j \\ (1-p_d)/\eta & \text{if} \quad l = k \\ (1-p_d)/(\eta(J_1-2)) & \text{if} \quad l \neq j, k \end{cases} , \qquad (11)$$

for $l \in \{1, .., J_1\}$ and $\eta > 1$. The idea here is that when we choose η close to 1, the category with the highest frequency has a relatively high probability to

Table 4: *Maximum of μ for values of (M, D, E) with frequency 1 in sample s_3 when using fine-tuning for all three variables.*

	p_d				
	0.95	0.90	0.85	0.80	0.70
no fine-tuning	0.93	0.83	0.69	0.55	0.31
fine-tuning 1 column where $\eta = 1.001$	0.92	0.80	0.66	0.51	0.28
fine-tuning 2 columns where $\eta = 1.001$	0.91	0.74	0.57	0.42	0.20
fine-tuning 3 columns where $\eta = 1.001$	0.90	0.72	0.52	0.34	0.15

change into the category with the lowest frequency. Assuming a link between an unsafe combination and a low frequency in the original sample, this idea explicitly supports the concept of PRAM: an unsafe combination c is after PRAM protected by creating new combinations c from combinations that have high frequencies in the original sample.

In the same way other columns in P_{X_1} can be fine-tuned. For example, the second column chosen is the column that corresponds to the category of X_1 with the second highest frequency in sample s, and the chosen row is now the row that corresponds to the category of X_1 with the second lowest frequency in sample s.

Table 4 presents results for sample s_3 when the transition matrices of M, D, and E are fine-tuned. The advantage of fine-tuning the transition matrices is dependent of the data and on the size of the sample. One can see that the idea works, e.g., if $p_d = 0.80$, fine-tuning can decrease the maximum of μ from 0.55 to 0.34.

Even after using fine-tuning, the maximum of μ is still quite large. Additional simulations, not reported, show that the maximum of μ decreases rapidly when sample size is increased. Conclusion and advice: determine a largest tolerated μ and check all combinations of three identifying variables and use fine-tuning. The protection offered by PRAM dependents on p_d, but also very much on the sample size.

Table 5: *Frequencies before PRAM and standard errors of estimated frequencies after PRAM of the variable M in sample s_3.*

	p_d					
	0.95		0.90		0.85	
$\boldsymbol{f}$	standard error of $\hat{\boldsymbol{f}}$ using P_M	$\overleftarrow{P}_M$	P_M	$\overleftarrow{P}_M$	P_M	$\overleftarrow{P}_M$
99	5.27	4.08	7.87	4.77	10.23	4.87
378	6.33	5.62	9.40	7.35	12.13	8.30
353	6.24	5.52	9.27	7.21	11.97	8.11
471	6.65	5.95	9.85	7.83	12.69	8.90
525	6.83	6.12	10.11	8.08	13.01	9.21
551	6.78	6.08	10.04	8.02	12.93	9.13
169	5.55	4.64	8.28	5.77	10.74	6.20

8.3 Information loss in frequency estimation

To investigate information loss due to PRAM, this section discusses an example with univariate frequency estimation with respect to the variable M and bivariate frequency estimation with respect to the variables S and E in sample s_3. We illustrate the difference between using P_X and $\overleftarrow{P}_X$ by comparing standard errors in estimating the univariate frequencies of variable M. In the following we assume that the released sample frequencies of M to be equal to the expected released sample frequencies. That is, released sample frequencies $\boldsymbol{f}^*$ are given by $\mathbb{E}(\boldsymbol{F}^*|\boldsymbol{f}) = P_M\boldsymbol{f}$, where $\boldsymbol{F}^*$ and $\boldsymbol{f}$ are defined with respect to M. It follows that in this situation $\hat{\boldsymbol{f}} = \boldsymbol{f}$, so that (5) can be used to compare variances. To estimate the standard errors when using calibration, we use $\overleftarrow{P}_M$ in (5) instead of P_M^{-1}. Table 5 presents standard errors of estimated frequencies for different choices of p_d. The example shows that $\overleftarrow{P}_M$ is more efficient than P_M, a difference that becomes more striking when p_d is smaller.

In the bivariate situation calibration probabilities and proportions do not always work well. To illustrate this, the following example is about frequency estimation of

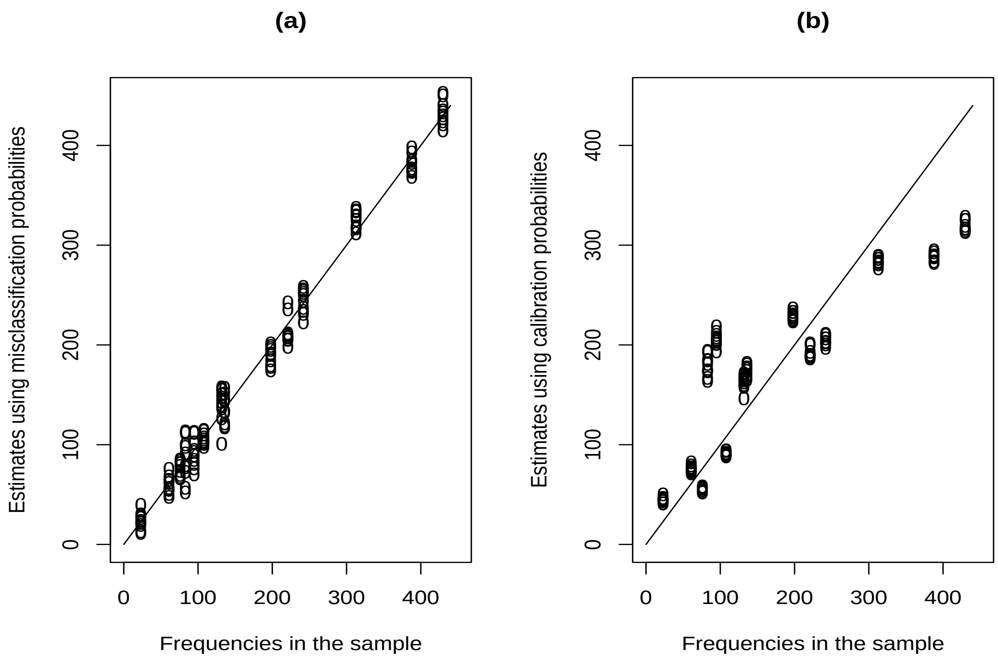

Figure 1: *Estimating frequencies of $X = (S, E)$ after applying PRAM to S and E in sample s_3 with p_d=0.85 in 10 simulations.* (a) *Using misclassification probabilities.* (b) *Using calibration probabilities.*

variable $X = (S, E)$ that has 14 categories. The χ^2-test of independence between S and E yields 529.55, where df = 6 and the p-value < 0.00. It is this lack of independence between the variables that causes calibration probabilities to perform badly. PRAM was applied 10 times to both S and E with $p_d = 0.85$. Figure 1 shows the estimation of the frequencies of X using X^* and $P_S \otimes P_E$ versus using X^* and $\overleftarrow{P}_S \otimes \overleftarrow{P}_E$. From the figure it is clear that the misclassification probabilities perform better, i.e., the points $(f_j, \hat{f}_j)$ are closer to the identity line then the points $(f_j, \tilde{f}_j)$, $j \in \{1, ..., 14\}$. The variance is less when $\overleftarrow{P}_S \otimes \overleftarrow{P}_E$ is used, but the figure shows that in that case estimates are biased. Violating the independence assumption regarding the use of $\overleftarrow{P}_S \otimes \overleftarrow{P}_E$ has sever consequences.

Table 6: *Actual coverage percentage w.r.t.* $X = (S, E)$ *for sample* s_3 *and 1000 simulated perturbed samples, where* $p_d = 0.85$.

	ACP using			ACP using	
category	$P_S \otimes P_E$	$P_S^\circ \otimes P_E^\circ$	category	$P_S \otimes P_E$	$P_S^\circ \otimes P_E^\circ$
(1,1)	94.8	98.2	(2,1)	94.5	97.8
(1,2)	95.3	98.1	(2,2)	95.6	97.7
(1,3)	95.7	97.7	(2,3)	95.0	97.9
(1,4)	95.4	98.0	(2,4)	95.7	98.8
(1,5)	95.5	99.1	(2,5)	93.5	97.3
(1,6)	95.1	97.4	(2,6)	93.8	97.5
(1,7)	96.0	98.3	(2,7)	96.0	98.7

Misclassification proportions are close to misclassification probabilities in the above example. Compare for instance

$$P_S = \begin{pmatrix} 0.85 & 0.15 \\ 0.15 & 0.15 \end{pmatrix} \quad \text{and} \quad P_S^\circ = \begin{pmatrix} 0.854 & 0.154 \\ 0.156 & 0.156 \end{pmatrix}.$$

A simulation study can be used to investigate the performance of misclassification probabilities versus misclassification proportions. The study compares using $P_S \otimes P_E$ versus using $P_S^\circ \otimes P_E^\circ$ by looking at the actual coverage percentage (ACP), which is the percentage of the replicated perturbed samples for which the confidence interval of the estimated frequency covers the actual frequency in the original sample. We used sample s_3, $p_d = 0.85$ and 1000 simulated perturbed samples. Table 6 shows that misclassification proportions perform better than then misclassification probabilities. The mean value of ACP when using $P_S \otimes P_E$ equals 95.14, and the mean value of ACP when using $P_S^\circ \otimes P_E^\circ$ equals 98.04. (A Paired t-Test yields a p-value < 0.00.) So, although the transition matrices are quite alike at first sight, misclassification proportions perform best.

9 Conclusion

The paper shows that the analysis of PRAM data is more efficient when misclassification proportions are used instead of misclassification probabilities. Calibration probabilities and calibration proportions work fine in the univariate case, but cause serious bias in the multivariate case. Since in most situations the user of PRAM data will be interested in multivariate analysis, it seems wise not to release calibration probabilities or calibration proportions along with the PRAM data. The two measures for disclosure risk that are used in this paper show that PRAM helps in protecting the identity of respondents.

Given that releasing misclassification proportions makes PRAM more efficient with respect to information loss, it is still an open question how this works out when PRAM is compared to other SDC methods, see Domingo-Ferrer and Torra (2001). It might be worthwhile to state that PRAM was never meant to replace existing SDC methods. Working with PRAM data and taking into account the information about the misclassification in the analysis might be quite a burden for some researchers. However, when researchers are interested in specific details in data, details that might disappear when, e.g., global recoding is used, PRAM can be a solution. Note that PRAM is statistically sound. Data are perturbed, but information about the perturbation can be used. Although estimates will have extra variance due to the perturbation, they will be unbiased.

Since the misclassification proportions provide more information about the original sample than the misclassification probabilities, one should consider the question whether providing these proportions increases the disclosure risk. Since the privacy protection that is offered by PRAM is at the record level, we do not think that disclosure risk increases when misclassification proportions are released. With these proportions, sample frequencies of the identifying variables can be deduced, but these frequencies are not sensitive information. Note also that when one works with the measures for disclosure risk discussed in Section 6, the risk does not change when misclassification proportions are released.

Appendix A

The following shows that $\boldsymbol{f} = \overleftarrow{P}_X \mathbb{E}[\boldsymbol{F}^*|\boldsymbol{f}]$. First note that $\mathbb{E}[\boldsymbol{F}^*|\boldsymbol{f}] = P_X \boldsymbol{f}$ and that entries $\overleftarrow{p}_{jk}$ of $\overleftarrow{P}_X$ are defined as $\overleftarrow{p}_{jk} = (p_{kj} f_j)(\sum_{j_0=1}^{J} p_{kj_0} f_{j_0})^{-1}$ for $k, j \in \{1, ..., J\}$. For each $j \in \{1, ..., J\}$ we have

$$\begin{aligned}\Big(\overleftarrow{P}_X \mathbb{E}[\boldsymbol{F}^*|\boldsymbol{f}]\Big)(j) &= \sum_{k=1}^{J} \overleftarrow{p}_{jk} \Big(\mathbb{E}[\boldsymbol{F}^*|\boldsymbol{f}]\Big)(k) \\ &= \sum_{k=1}^{J} \overleftarrow{p}_{jk} \Big(\sum_{j_0=1}^{J} p_{kj_0} f_{j_0} \Big) = \sum_{k=1}^{J} p_{kj} f_j = f_j,\end{aligned}$$

since the columns of P_X sum up to one. So $\boldsymbol{f} = \overleftarrow{P}_X P_X \boldsymbol{f}$ and $\boldsymbol{f}$ is an eigenvector of $\overleftarrow{P}_X P_X$ with eigenvalue 1.

In general, $\overleftarrow{P}_X \neq P_X^{-1}$. To illustrate this, let $R = \overleftarrow{P}_X P_X$. The entries of R are $r_{ij} = \sum_{k=1}^{J} \overleftarrow{p}_{ik} p_{kj}$, for $i, j \in \{1, ..., J\}$. Assume that the entries of P_X are all > 0 and that $\boldsymbol{f}_j > 0$, for $j \in \{1, ..., J\}$. Then $\overleftarrow{p}_{jk} > 0$, for $j, k \in \{1, ..., J\}$ and $r_{ij} > 0$, for $i, j \in \{1, ..., J\}$. In this case, R is not the identity matrix and consequently $\overleftarrow{P}_X \neq P_X^{-1}$. A more intuitive explanation is that $\overleftarrow{P}_X$ changes when the survey data change, whereas P_X can be determined independently from the data and hence does not necessarily change when the data change. Therefore, it is always possible to cause $\overleftarrow{P}_X \neq P_X^{-1}$ by changing the data.

Appendix B

The following derives the maximum likelihood properties of (3) and (8). The reasoning is the same as in Hochberg (1977), but simpler, since in the PRAM situation calibration probabilities do not have to be estimated. Also, we show that the reasoning applies both to (3) and to (8).

Assume that $X_1, ...X_n$ are independently multinomially distributed with parameter vector $\boldsymbol{\pi} = (\pi_1, .., \pi_J)^t$, where $\pi_j > 0$ for $j \in \{1, .., J\}$, and $\sum_{j=1}^{J} \pi_j = 1$. Consider the transformation $\boldsymbol{\pi}^* = P\boldsymbol{\pi}$, where P is a $J \times J$ transition matrix, i.e., columns sum up to one and $p_{kj} \geq 0$ for $k, j \in \{1, .., J\}$. Assume that P is nonsingular. Let the distribution of X^* be given by $\mathbb{P}(X^* = k) = \pi_k^*$, for $k \in \{1, .., J\}$. It follows that $X_1^*, X_2^*, ..., X_n^*$ are multinomially distributed with parameters n and $\boldsymbol{\pi}^*$. Indeed, $\pi_k^* = p_{k1}\pi_1 + ... + p_{kJ}\pi_J > 0$ for $k \in \{1, .., J\}$ and

$$\sum_{k=1}^{J} \pi_k^* = \left(\sum_{l=1}^{J} p_{l1}\right)\pi_1 + ... + \left(\sum_{l=1}^{J} p_{lJ}\right)\pi_J = 1.$$

The likelihood L^* for $\boldsymbol{\pi}^*$ and observed $\boldsymbol{x}^* = (x_1^*, x_2^*, ..., x_n^*)^t$ is well known. Let $\boldsymbol{f}^* = (f_1^*, f_2^*, ..., f_J^*)^t$ denote the observed cell frequencies. The MLE is given by $\hat{\boldsymbol{\pi}}^* = \boldsymbol{f}^*/n$ and has covariance matrix $\Omega = [\text{Diag}(\boldsymbol{\pi}^*) - \boldsymbol{\pi}^*(\boldsymbol{\pi}^*)^t]/n$, where $\text{Diag}(\boldsymbol{\pi}^*)$ is the diagonal matrix with the diagonal entries given by the elements of $\boldsymbol{\pi}^*$.

Next we can use the invariance property of maximum likelihood. Define the transformation $g(\boldsymbol{\pi}^*) = P^{-1}\boldsymbol{\pi}^*$. Since g is one-to-one, it follows from $L^*(\boldsymbol{\pi}^*|\boldsymbol{x}^*)$ and $\boldsymbol{\pi} = g(\boldsymbol{\pi}^*)$ that the likelihood for $\boldsymbol{\pi}$ is given by $L^*(g^{-1}(\boldsymbol{\pi})|\boldsymbol{x}^*)$ which is maximized for $\hat{\boldsymbol{\pi}} = g(\hat{\boldsymbol{\pi}}^*) = P^{-1}\hat{\boldsymbol{\pi}}^*$. Consequently, when $\hat{\boldsymbol{\pi}} \in (0, 1)^J$, it is the MLE. Since g has a first order derivative, the covariance matrix of $\hat{\boldsymbol{\pi}}$ can be obtained using the delta-method, see, e.g., Agresti (1990, Chapter 12). We have $\partial g(\boldsymbol{\pi}^*)/\partial \boldsymbol{\pi}^* = (P^{-1})^t$ and the covariance matrix of $\hat{\boldsymbol{\pi}}$ is given by $n^{-1}P^{-1}\Omega(P^{-1})^t$.

So, with respect to (3) maximum likelihood properties are proven by taking $P = P_X$ and obtaining $\hat{\boldsymbol{\pi}} = g(\hat{\boldsymbol{\pi}}^*) = P_X^{-1}\hat{\boldsymbol{\pi}}^*$. With respect to (8), the misclassification design is described by $\boldsymbol{\pi} = \overleftarrow{P}_X \boldsymbol{\pi}^*$, so $P = \overleftarrow{P}_X^{-1}$ and the MLE is given by $\tilde{\boldsymbol{\pi}} = g(\hat{\boldsymbol{\pi}}^*) = \overleftarrow{P}_X \hat{\boldsymbol{\pi}}^*$.

Appendix C

Let P°_{kj} denote the stochastic variable of the kj-th entry of P°_{X} and C_{kj} the stochastic variable of the kj-th cell in the cross-classification X^* by X. It follows that C_{kj} has a binomial distribution with parameters f_j and p_{kj}. Consequently, $\mathbb{E}[P^{\circ}_{kj}|\boldsymbol{f}] = \mathbb{E}[C_{kj}/f_j|\boldsymbol{f}] = f_j p_{kj}/f_j = p_{kj}$ and in expectation P°_{X} equals P_X. Since $C_{k_1 j_1}$ and $C_{k_2 j_2}$ are independent given $\boldsymbol{f}$, it follows that $\mathbb{E}[P^{\circ}_{k_1 j_1} P^{\circ}_{k_2 j_2}|\boldsymbol{f}] = p_{k_1 j_1} p_{k_2 j_2}$. So, in expectation, $P^{\circ}_{X_1} \otimes P^{\circ}_{X}$ equals $P_{X_1} \otimes P_{X_2}$.

We define $\overleftarrow{P}^{\circ}_{jk} = C_{kj}/(F^*_k + \varepsilon)$ where ε is a small positive value. Using the delta method, see, e.g., Rice (1995, Section 4.6), we obtain

$$\mathbb{E}[\overleftarrow{P}^{\circ}_{jk}|\boldsymbol{f}] \approx \frac{\mathbb{E}[C_{kj}|\boldsymbol{f}]}{\mathbb{E}[F^*_k|\boldsymbol{f}]} + \frac{1}{\mathbb{E}[F^*_k|\boldsymbol{f}]^2}\left(V[F^*_k|\boldsymbol{f}]\frac{\mathbb{E}[C_{kj}|\boldsymbol{f}]}{\mathbb{E}[F^*_k|\boldsymbol{f}]} - \rho\sqrt{V[C_{kj}|\boldsymbol{f}]V[F^*_k|\boldsymbol{f}]}\right),$$

where $\mathbb{E}[C_{kj}|\boldsymbol{f}] = f_j p_{kj}$ and ρ is the correlation between C_{kj} and F^*_k. From this we see that the difference between $\mathbb{E}[\overleftarrow{P}^{\circ}_{jk}|\boldsymbol{f}]$ and $\overleftarrow{p}_{jk}$ will be small when $V[F^*_k|\boldsymbol{f}]$ is small and $\mathbb{E}[F^*_k|\boldsymbol{f}]$ is large.

References

Agresti, A. (1990). Categorical Data Analysis. New York: Wiley.

Bethlehem J.G., Keller, W.J. and Pannekoek, J. (1990). Disclosure Control of Microdata. Journal of the American Statistical Association, 85, 38-45.

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

De Wolf, P.-P., Gouweleeuw, J.M., Kooiman, P., and Willenborg, L.C.R.J (1997). Reflections on PRAM. Research paper no. 9742, Voorburg/ Heerlen: Statistics Netherlands.

Domingo-Ferrer, J., and Torra, V. (2001). A Quantitative Comparison of Disclosure Control Methods for Microdata. In Confidentiality, Disclosure, and Data Access, Theory and Practical Application for Statistical Agencies. Eds. Doyle, P., Lane, J., Zayatz, L., and Theeuwes, J., Amsterdam: North-Holland.

Elamir, E.A.H., and Skinner, C.J. (2003). Disclosure Risk Assessment for Microdata, the Treatment of Measurement Error. Manuscript.

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. Journal of Official Statistics, 14, 463-478.

Hochberg, Y. (1977). On the Use of Double Sampling Schemes in Analyzing Categorical Data with Misclassification Errors. Journal of the American Statistical Association, 72, 914-921.

Kooiman, P., Willenborg, L.C.R.J., and Gouweleeuw, J.M. (1997). PRAM, a Method for Disclosure Limitation of Microdata. Research paper no. 9705, Voorburg/Heerlen: Statistics Netherlands.

Kuha, J.T., and Skinner, C.J. (1997). Categorical Data Analysis and Misclassification. In Survey Measurement and Process Quality, Eds. Lyberg, L., Biemer, P., Collins, M, De Leeuw, E, Dippo, C, Schwarz, N, and Trewin, D., New York: Wiley.

Rice, J.A. (1995). Mathematical Statistics and Data Analysis. Belmont: Duxbury Press.

Rosenberg, M.J. (1979). Multivariate Analysis by a Randomised Response Technique for Statistical Disclosure Control. Ph.D. dissertation, University of Michigan, Ann Arbor.

Skinner, C.J., and Elliot, M.J. (2002). A Measure of Disclosure Risk for Microdata. Journal of the Royal Statistical Society, B, 855-867.

Van den Hout, A., and Van der Heijden, P.G.M. (2002). Randomized Response, Statistical Disclosure control and Misclassification, a Review. International Statistical Review, 70, 269-288.

Warner, S.L. (1965). Randomized Response: a Survey Technique for Eliminating Answer Bias. Journal of the American Statistical Association, 60, 63-69.

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