WP. 11
ENGLISH ONLY

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

EUROPEAN COMMISSION STATISTICAL OFFICE OF THE EUROPEAN COMMUNITIES (EUROSTAT)

Joint UNECE/Eurostat work session on statistical data confidentiality
(Geneva, Switzerland, 9-11 November 2005)

Topic (ii): Disclosure risk, information loss and usability of data

ROMM METHODOLOGY FOR MICRODATA RELEASE

Invited Paper

Submitted by Carnegie Mellon University, United States and
Universidad de Alicante, Spain [1]

[1] Prepared by Daniel Ting and Stephen Fienberg, Carnegie Mellon University (fienberg; dting@stat.cmu.edu) and Mario Trottini, Universidad de Alicante (mario.trottini@ua.es).

ROMM Methodology for Microdata Release

Daniel Ting*, Stephen Fienberg**, Mario Trottini***

* Department of Statistics, Carnegie Mellon University, Pittsburgh, PA 15213, USA. (dting@stat.cmu.edu)

** Department of Statistics, Center for Automated Learning and Discovery, Cylab, Carnegie Mellon University, Pittsburgh, PA 15213, USA. (fienberg@stat.cmu.edu)

*** Departmento de Estadística e I.O , Universidad de Alicante, Spain (mario.trottini@ua.es)

Abstract. Statistically defensible methods for disclosure limitation allow data users to make inferences about parameters in a model similar to those that would be possible using the original unreleased data. We present a new perturbation method for protecting confidential attributes in continuous microdata—**R**andom **O**rthogonal **M**atrix **M**asking (ROMM) which preserves the sufficient statistics for multivariate normal distributions, and thus is statistically defensible. ROMM encompasses all methods that preserve these statistics and can be restricted to provide "small" perturbations. We discuss methods for evaluating the disclosure risk and data utility of ROMM.

1 Introduction

Statistical agencies and publicly-funded researchers are under a dual obligation to share data with others, especially in the form of detailed microdata, and at the same time preserving the confidentiality of the respondents who provided these data. To protect the data, they typically must do something beyond removing obvious identifiers from the individual records. When the data come from a sample survey, the sampling rate may be sufficient to provide suitable protection, although not necessarily for local geographic areas. The natural question is then: "How should we modify the data?" To answer this question we need to ask: "What algorithm should one used to modify the data?" "How can we ascertain the extent to which confidentiality is protected?" "How useful are the modified data?"

For continuous microdata, among the most popular methods has been the addition of noise, e.g., see Kim (1986) and Kim and Winkler (1995), but other proposals include Muralidhar, et al. (1999), "statistical obfuscation" (Burridge (2003)), PRAM (Gouweleeuw, et al. (1998)), rank-swapping (Reiss, et al. (1984)) and multiple imputation (Rubin (1993)). These techniques succeed at some level in protecting the confidentiality of the data, although there is disagreement as to the extent of the protection, but the degree to which they provide useful data for the analyst is a far more contentious issue. Fienberg (2005), Rubin (1993), and others have argued that statistically defensible methods for disclosure limitation need to allow the data analyst to make inferences about parameters of interest in a model applicable to the original unreleased data. If the data are transformed then one way to achieve this is to provide details of the method to allow the creation of a usable likelihood function for the true unreleased data, as with PRAM. Another strategy is to preserve essential features of the data as part of the transformation process. Burridge (2003) does this for continuous microdata by preserving the minimal sufficient statistics of the data under an assumption of multivariate normality, i.e., the mean and covariance matrix. Such a choice means that the user who sets out to apply a normal-distribution-based multivariate method will get the same estimates for the underlying parameters. But when such estimates are applied to the transformed data, to produce residuals for example, the new results should reflect the added uncertainty associated with the perturbation process.

In Section 2, we propose a new perturbation method, *Random Orthogonal Matrix Masking* (ROMM), for continuous microdata. ROMM's principal features are:

- It preserves *sample* means and *sample* covariances, the sufficient statistics of a multivariate normal, and hence it exactly preserves linear regression estimates.
- It controls the magnitude of the perturbation, so useful analyses can often be performed on the perturbed data even when the underlying model is not assumed to be a multivariate normal.

We describe the implementation of the method in detail in Section 3, and then, in Section 4, we discuss evaluating the level of confidentiality protection and the usefulness of the perturbed data formally in terms of risk and utility.

2 Random Orthogonal Matrix Masking

We begin by introducing **R**andom **O**rthogonal **M**atrix **M**asking (ROMM). We then link ROMM to matrix masking as described by Duncan and Pearson (1991) and we show that it encompasses other previously proposed methods.

The procedure for ROMM is as follows:

1. Remove identifying variables such as name, address, and social security other forms of publicly accessible identification numbers and represent the resulting data as an $n \times k$ matrix x.
2. Generate a random orthogonal matrix, t, from a distribution G defined on the group of $n \times n$ orthogonal matrices which keep 1_n invariant, i.e., $t1_n = 1_n$ where 1_n is the column vector consisting of n 1's.
3. Apply the orthogonal operator, t to the original data x to produce perturbed microdata y:
$$y = tx.$$
4. Release to the users:
   (a) The output of the transformation, y;
   (b) The information that y has been obtained applying to the original data an orthogonal operator randomly generated from a distribution G;
   (c) The exact distribution G.

ROMM is a specific case of matrix masking, which for the $n \times k$ data matrix x involves the transformation
$$x \longrightarrow y = AxB + C.$$
For ROMM, B is the identity, C is the zero matrix, and the class of masks consists of those A's that are random orthogonal matrices drawn from some known distribution G.

ROMM was designed so that the sample means and sample covariance matrix for x and y are the same. Furthermore, we can show that, for *any* pair of matrices x', y' with the sample means and sample covariance matrix, there exists an orthogonal matrix t that keeps 1_n invariant and $y' = tx'$. Because the method of "statistical obfuscation" proposed by Burridge (2003) preserves the sample means and sample covariance matrix, it is a special case of ROMM. We give two theorems that formally codify these statements and refer the interested reader to the proofs in Ting, et al. (2005).

Theorem 1 *Let $\bar{x}$ and Σ_x be the sample mean and the sample covariance matrix of the original microdata and let $\bar{y}$ and Σ_y be the corresponding quantities in the masked microdata produced by ROMM. Then*

$$\bar{x} = \bar{y} \quad \textit{and} \quad \Sigma_x = \Sigma_y.$$

Theorem 2 *Let M be any data masking procedure that generates a random microdata, y, with the same sample mean and sample covariance matrix as the original microdata. Then M is a special case of ROMM for a suitable choice of the "parameter" G.*

The preservation of the sample means and sample covariance matrix has special practical and theoretical features. On practical level, simple linear regression estimates are preserved exactly. On a theoretical level, the sufficient statistics are preserved when the underlying distribution is assumed to be multivariate normal.

The second feature of ROMM, namely the ability to control the magnitude of perturbation, is achieved through an appropriate choice of distribution G. A perturbation is considered small if the (Riemannian) distance d of the orthogonal matrix to the identity is close to zero. Section 3 describes some choices of G for which draws from G tend to correspond to small perturbations.

2.1 Comparison With Additive Perturbation Methods

Much of the analysis of additive perturbation methods has focused on the effects on regression estimates and estimating the covariance matrix of the underlying distribution. How does ROMM compare with such methods in the regression setting?

Under the naive additive perturbation method, uncorrelated additive noise, the problem of estimating regression coefficients reduces to the well-known problem of estimatation with measurement error in the covariates. The usual regression estimates are biased towards 0 and inconsistent. Hence, they must be bias corrected. The resulting bias corrected estimates on perturbed data have greater variance than estimates based on unperturbed data. See Lechner and Pohlmeier (2004) for related and more general discussion.

With perturbed data obtained through bias corrected and correlated additive noise in Kim (1986) or with GADP in Muralidhar, et al. (1999), the sample mean and sample covariance estimates on the perturbed data are unbiased estimates of the true *estimates* obtained from the original data. While this means they are also unbiased estimates of the true parameters in a multivariate normal model, it underscores the fact that additional variability is introduced. The Rao-Blackwell theorem shows that this increase in variance is strict unless the sufficient statistics of the unperturbed data can be recovered exactly from the perturbed data. This is clearly not the case under these additive noise methods. Furthermore, the usual regression estimates under these additive noise methods are not necessarily unbiased. They are, however, consistent.

3 Implementation: Distributions on Orthogonal Operators

In this section we describe some choices for G, the distribution on orthogonal matrices, and show how an appropriate choice of the parameters for G results in small perturbations. To simplify the descriptions, we temporarily ignore the requirement that the vector 1_n must be held invariant. This deficiency is easily corrected by treating an $(n-1) \times (n-1)$ orthogonal matrix as a linear operator on the space orthogonal to 1_n and then extending it to hold 1_n invariant to obtain an $n \times n$ orthogonal matrix. We give the details of each sampling method and some additional distributions in (Ting, et al. (2005)).

3.1 Coordinate by Coordinate and Uniform Distributions

Using the idea that small perturbations are matrices "close to" the identity, we add a small amount of noise to the identity and then modify this matrix to be orthogonal. The algorithm is as follows.

1. Choose a parameter $\lambda > 0$ corresponding to the magnitude of perturbation.
2. Draw an $n \times n$ random matrix M with entries from a standard normal.
3. Put $P = I + \lambda M$.
4. Apply Gram-Schmidt and normalize the columns of P to obtain an orthonormal matrix T.

It is easy to see that, when $\lambda = 0$, T is the identity and no perturbation has occured. When $\lambda = \infty$, T is a draw from the uniform distribution on orthogonal matrices (according to Haar Measure) (see Eaton (1983, p. 234)).

3.2 Block Diagonal Distribution

Another approach to control the magnitude of perturbations is to consider orthogonal matrices with eigenvalues close to 1. For simplicity assume n is even. To sample from the block diagonal distribution, we perform the following steps.

1. Draw B from the uniform distribution on orthogonal matrices.
2. For $j = 1...n/2$, independently draw $b_j \sim Beta(\alpha, \beta)$ for some choice of parameters α and β. Put $\theta_j = 2\pi b_j - \pi$.
3. Let L be the block diagonal matrix where each block is the 2×2 matrix
$$\begin{pmatrix} cos\theta_j & -sin\theta_j \\ sin\theta_j & cos\theta_j \end{pmatrix}$$
4. Put $T = BLB^T$. T is then an orthogonal matrix with eigenvalues $e^{\pm i\theta_j}$. Furthermore, the support of the resulting block diagonal distribution contains all orthogonal matrices except for a set of measure 0.

It is easy to see that as $\alpha = \beta \to \infty$, the distribution of each b_j converges to 0, and T converges to I. Unlike the coordinate by coordinate distribution, however, the block diagonal distribution cannot reproduce the uniform distribution because the eigenvalues of orthogonal matrices drawn from the uniform distribution are highly correlated (see Diaconis and Shahshahani (1994)), and this distribution assumes independence of eigenvalues.

3.3 Example: Boston Housing Data

We demonstrate ROMM's performance using a subset of 4 variables from the Boston house price data in Harrison and Rubinfeld (1978):

Variable	Meaning
RM	average number of rooms per dwelling
PTRATIO	pupil-teacher ratio by town
LSTAT	% lower status of the population
MEDV	Median value of owner-occupied homes in $1000's

We may treat LSTAT as the sensitive variable, i.e., homeowners in a particular tract do not want to disclose what percentage of people in that tract have low socioeconomic status. The regression model of interest has MEDV as the dependent variable and the remaining ones as predictors. We drew the following random subset of 13 observations:

Obs	RM	PTRATIO	LSTAT	MEDV	Obs	RM	PTRATIO	LSTAT	MEDV
1	6.630	18.5	6.53	26.6	8	6.315	16.6	7.60	22.3
2	5.986	19.1	14.81	21.4	9	6.023	18.4	11.72	19.4
3	5.709	14.7	15.79	19.4	10	6.251	20.2	14.19	19.9
4	5.877	14.7	12.14	23.8	11	5.757	20.2	10.11	15.0
5	6.402	14.7	11.32	22.3	12	5.304	20.2	26.64	10.4
6	6.782	15.2	6.68	32.0	13	6.425	20.2	12.03	16.1
7	6.433	19.1	9.52	24.5					

We give two perturbed datasets, one using ROMM and one using the bias corrected, correlated additive noise method described in Kim (1986). The weight of the noise in the additive noise method is $\sqrt{c} = 1/2$. The ROMM pertubation distribution is the coordinate-by-coordinate distribution with $\lambda = 1/3$. The difference of the perturbed data under each method and the original data is given below. It is evident that that magnitude of the differences is in general much larger under ROMM.

ROMM				Additive Noise			
RM	PTRATIO	LSTAT	MEDV	RM	PTRATIO	LSTAT	MEDV
-0.253	3.725	4.721	-7.681	-0.337	1.509	1.848	-2.173
-0.529	2.006	3.843	-9.182	0.182	2.249	-1.479	-1.238
0.494	0.738	-12.337	5.930	-0.036	2.335	-0.761	-3.027
0.045	0.880	1.249	-0.095	0.235	-0.132	-3.451	0.908
0.223	1.313	-1.086	3.740	-0.057	1.274	3.057	-3.099
-0.183	1.841	0.338	-3.079	-0.324	0.380	1.999	-3.797
-0.269	2.093	1.883	-4.906	-0.281	-1.957	3.303	0.980
0.139	-0.867	0.348	3.018	0.175	-0.046	-0.831	0.765
0.414	0.367	-1.688	0.357	0.053	-1.423	-2.049	1.165
0.212	-3.778	-5.139	4.602	0.086	-0.116	0.367	0.883
-0.437	-3.218	10.437	-2.359	-0.107	0.195	2.658	0.143
0.601	-2.851	-6.872	12.014	0.390	-0.581	-8.411	4.243
-0.457	-2.249	4.302	-2.361	-0.468	0.254	6.544	-5.830

The following sample variances give a rough idea of how much greater the magnitudes of the differences are for ROMM:

	RM	PTRATIO	LSTAT	MEDV
ROMM	0.145	5.618	33.307	34.327
Additive noise	0.066	1.641	13.950	7.346

The linear regression estimates for the original data set (and ROMM) and for the additive noise version are as follows:

Original Data (and ROMM)			Additive Noise		
Variable	Estimate	Std. Error	Variable	Estimate	Std. Error
(Intercept)	-5.5641	23.6517	(Intercept)	-2.4696	25.1829
RM	7.4488	3.3663	RM	6.2655	3.5843
PTRATIO	-0.9557	0.3691	PTRATIO	-0.3930	0.4329
LSTAT	-0.1770	0.2741	LSTAT	-0.6780	0.2995

The differences between regression estimates are substantial in this case. For `PTRATIO` and `LSTAT` the difference in coefficients is greater than the estimated standard error for each coefficient. Further, we note that the inferences under the additive noise method do not reflect the added uncertainty of the perturbation.

4 Disclosure Risk and Data Utility

Like any data masking procedure aimed at finding a suitable balance between *safety* and *usability* of the perturbed data, ROMM relies on an implicit assumption that targets of potential intruders do not overlap with targets of legitimate data users, c.f. Trottini (2004). The underlying assumption is that researchers, policy makers, and public opinion are interested in statistical analyses aimed at discovering and making inferences about general features of the population represented by the data (e.g., association among variables, or the models for different type of phenomena) while intruders are interested in identifying confidential information about individual respondents. Within this framework, for any given unperturbed data set the knowledge of the distribution that has generated the data is "sufficient" for any statistical analysis that legitimate data users might wish to perform. ROMM exploits this idea by providing a perturbation method that preserves general features of the distribution while increasing the difficulty for an intruder to recover confidential information about individual respondents.

A rigorous assessment of disclosure risk and utility requires:

- A model for *users'* behaviors when the ouput of ROMM is released,
- An assessment of agency uncertainty about this model's inputs (*users'* targets, prior information, estimation procedure, etc.)
- A formalization of agency's perception of the consequences of data users' actions and of agency's preference structure for consequences of *users'* actions (see Trottini (2004)).

Because of space limitations, here we consider a simplified scenario where: (i) the modeling of *users* and agency's behaviors does not take explicitly into account some relevant aspects of the problem, such as agency's perception of usefulness in terms of model checking, diagnostics, and feasibility of the users' inferences under the released data, and (ii) the agency has no uncertainty about the *users*' model inputs. We can extend the results to more realistic scenarios by explicitly incorporating agency's uncertainty on *users'* model inputs as described in Trottini (2004), and by explicitly formalizing agency's perception of usefulness in terms of model checking, diagnostics, and feasibility of users inferences using a suitable structuring of objectives and attributes (see Trottini (2005)).

4.1 Notation and Posterior Distribution

We first fix some notation, and, before evaluating risk and utility, we give a formula for the posterior distribution of the unreleased data, x. Then we use these in evaluating utility under non-normality assumptions and in evaluating disclosure risk.

1. Let X and T be random variables representing, respectively, the unperturbed data and the random orthogonal matrix used to transform the unperturbed data.

2. Let x_{orig} be the realized value of X data (i.e., the true values for the unperturbed data) and let t be the realized value for T.

3. Define $Y = TX$ to be the random variable representing the perturbed data. Then $y = tx_{orig}$ is the realization of Y.

4. Define E to be any external knowledge available to the user. In particular E may contain nonconfidential values in x_{orig} that may be used for record linkage or to undo the perturbation t.

5. Let $\tilde{m}$ be the masked data set produced by ROMM. Following Trottini (2004) we can represent $\tilde{m}$ as a pair $\tilde{m} = (y, I(T, G))$ where $I(T, G)$ represents the information provided by the agency to the *users* about the transformation (in this case, the information that the released microdata has been obtained by applying to the original data a random orthogonal operator, T, generated from a distribution G.)

6. Let $A(x, y) = \{t : y = tx\}$ be the set of orthogonal matrices such that 1_n is invariant under t and that take x to y.

The posterior distribution is then given by

$$\pi(x|\tilde{m}, E) \propto \pi(x|E) \cdot L(\tilde{m}; x, E) = \pi(x|E) \cdot \int_{A(x,y)\cap Supp(G)} dG(t)$$

where $Supp(G)$ is the support of G. Note that we assume G is given to all users since valid inferences are not possible otherwise. This formula also assumes that T is independent of X which may not be true. We give the full derivation of this result and a discussion of the case when T and X are dependent in Ting, et al. (2005). Another point worth noting is that, for $n \times k$ matrices x and y, the dimension of $A(x, y)$ is $(n-k)(n-k-1)/2$. Thus we need to compute a high-dimensional integral, and this is difficult. For $n = 100$ and $k = 10$ the dimension is 4000! Under some assumptions on the users' prior and on G, however, it is possible to sample from the posterior distribution (see Ting, et al. (2005)).

4.2 Data Utility

Utility Under Normality To assess data utility associated with ROMM, we first consider the procedure under normality assumptions. For this case, the notation given above is not important, but it will be when we consider the nonnormal case. If the original data are independent and identically distributed (i.i.d.) realizations from a multivariate normal distribution, the output of ROMM is a random sample from the the same multivariate normal distribution (any orthogonal transformation that preserves 1_n preserves the multivariate normal distribution of the original data). Thus, regardless of the distribution, G, used to generate the orthogonal operator and regardless of the inferences of interest for legitimate data users, under normality the ROMM procedure guarantees maximum data utility. The trade-off dilemma in this case is trivial. The statistical agency should choose the "noise parameter," G, to minimize the risk of disclosure, since data utility is constant (and maximum!) as a function of G.

Utility Under Non-normality When the data are not normal, ROMM preserves the mean vector and the covariance matrix of the unperturbed data but no longer the distribution. Heuristically, the idea of using small perturbations suggests that using y directly as input may still be useful in exploratory data analyses. For more rigorous analyses, however, legitimate data users interested in inferences other than the mean and covariances, e.g., in quantiles or mixture models, cannot directly use the output y as input to their standard statistical analyses and expect the resulting inferences to be valid. To make valid inferences, the output y is relevant only to the extent that provides information about the transformation, t, that has been applied to the original data and thus, indirectly, about the original data x_{orig}. Formally, a user must make inferences through the posterior distribution of x_{orig} given y and any additional information E the user has.

For a given legitimate data user's target, denote by $Z(x_{orig})$ the inferences that legitimate data users would make if they had access to the original data. His/her inference under the released ROMM data will be:

$$Z(\tilde{m}) = \int Z(x) \cdot \pi(x|\tilde{m}, E)dx \tag{1}$$

where $\pi(x|\tilde{m}, E)$ is the *user's* posterior distribution for the unperturbed data. Note, however, that the computation of $\pi(x|\tilde{m}, E)$ in this case can be very complex.

Depending on the agency's interpretation of *usefulness* of the data, we can define different measures of data utility as a function of (1). If, for example, the agency is concerned with minimizing the difference between data users inferences with the original data and the corresponding inferences with the masked data set, then a generic measure of data utility might be:

$$DU = D(Z(\tilde{m}), Z(x_{orig})),$$

where D is a distance metric that depends on the nature of the target Z and on the features of the inferences that should to be preserved. See Trottini (2004) for some examples.

4.3 Disclosure Risk

There are multiple ways in which a data intruder can pose a risk to a statistical agency in charge of releasing the microdata. The way we consider here is to take disclosure risk as the data intruder's utility. The problem then is similar to that of data utility for legitimate users. A data intruder has target $\zeta(i, x_{orig})$ where i corresponds to an individual with some row of values in x_{orig}. The disclosure risk may then be some distance metric between the target estimate and the true target value. Alternately, it be a reflection of the intruder's uncertainty of the target, such as the variance in the estimate of the target.

There is a subtle difference between this disclosure risk case and the case considered under the data utility section above. In this case, the row number r corresponding to the targetted individual is unknown to the intruder. This means that for data intruders, exchanging the rows of x_{orig} corresponds to changing the targetted individual when a particular row r is targetted. For a legitimate user's target, in many cases the rows corresponding to individuals may be exchanged without affecting the estimates. For example, the maximum of a list of numbers is invariant under permutation. Because of this difference, we must account for the uncertainty of r. To do this we write $x_{perm} = p^T x_{orig}$ where p^T is a permutation matrix that takes the r^{th} row to the 1^{st} row. The intruder's target $\zeta(i, x_{orig})$ then becomes a function $\zeta_{perm}(x_{perm})$ which does not directly depend on i. We may then treat p as the realization of a random variable P with some distribution on permutation matrices (not necessarily the uniform distribution). Since a permutation matrix is an orthogonal matrix that holds 1_n invariant, it follows that TP is an orthogonal matrix that holds 1_n invariant, and it has some distribution G_P. Since $y = (tp)x_{perm}$, the problem of estimating $\zeta(x_{perm})$ reduces to calculating the posterior as in the data utility case, but with G replaced by G_P. Note that when G is the uniform distribution, $G_P = G$ by the left invariance of G.

We provide more details in Ting, et al. (2005).

5 Conclusion

In this paper we introduced the disclosure limitation method of Random Orthogonal Matrix Masking (ROMM), and we demonstrated some of its theoretical properties. In particular, ROMM is designed to preserve the sufficient statistics of a multivariate normal distribution and thus preserves many common statistical quantities favored by users in their analyses, e.g., linear regression estimates. Furthermore, ROMM encompasses the entire class of perturbation methods that preserve these sufficient statistics. The specific procedures for ROMM we introduced suggest how it utilizes "small perturbations" of the data. We then consider the disclosure risk and data utility associated with ROMM. We showed how it may be assessed within a Bayesian framework using the posterior distribution of original data given the perturbed data and other additional information.

There is considerable work to be done to turn ROMM into a complete statistically-defensible method of disclosure limitation. For example, while the distributions on orthogonal matrices described in this paper heuristically favor small perturbations, we do not have theoretical results as to what the magnitude of the perturbation is, nor do we have theoretical results about the resulting distribution of perturbed data when the underlying distribution is not normal. From a practical perspective, although we gave a formula for the posterior, we also noted that it is difficult to calculate. In Ting, et al. (2005), we show how to sample from the posterior under certain assumptions about the prior. How these difficulties and assumptions will play out remains to be seen.

Acknowledgements

The research reported here was supported in part by NSF grant EIA-0131884 to the National Institute of Statistical Sciences, and by Army contract DAAD19-02-1-3-0389 to CyLab, at Carnegie Mellon University.

References

Burridge, J. (2003) Information Preserving Statistical Obfuscation, *Statistics and Computing*, 13, 4321–327.

Diaconis, P. and Shahshahani, M. (1994) On Eigenvalues of Random Matrices, *Journal of Applied Probability*, 10, 49–62.

Duncan, G.T., Pearson, R.W. (1991) Enhancing Access to Microdata While Protecting Confidentiality: Prospects for the Future, *Statistical Science*, 6, 219-232.

Eaton, M. (1983) *Multivariate Statistics: A Vector Space Approach*, Wiley, New York.

Fienberg, S. E. (2005). Allowing Access to Confidential Data: Some Recent Experiences and Statistical Approaches, Submitted for publication.

Gouweleeuw, J. M., Kooiman, P., Willenborg, L. C. R. J., and Wolf, P. P. de. (1998) Post Randomization for Statistical Disclosure Control: Theory and Implementation, *Journal of Official Statistics*, 14, 463–478.

Harrison, D. and Rubinfeld, D.L. (1978) Hedonic Prices and the Demand for Clean Air, *Journal of Environmental Economics and Management*, 5, 81–102.

Kim, J. J. (1986) A Method for Limiting Disclosure in Microdata Based on Random Noise and Transformation, *Proceedings of the Section on Survey Research Methods*, American Statistical Association, 303–308.

Kim, J. J., and Winkler, W. E. (1995) Masking Microdata Files, *Proceedings of the Section on Survey Research Methods*, American Statistical Association,114–119.

Lechner, S. and Pohlmeier, W. (2004) To Blank or Not to Blank? A Comparison of the Effects of Disclosure Limitation Methods on Nonlinear Regression Estimates. *Privacy in Statistical Databases 2004*, Springer-Verlag, 187–200.

Muralidhar, K., Parsa, R. and Sarathy, R. (1999) A General Additive Data Perturbation Method for Database Security, *Management Science*, 45, 1399–1415.

Reiss, S. P., Mark, J., and Dalenius, T. (1982) Non-reversible Privacy Transformations, *Proceedings of the ACM Symposium on Principles of Database Systems*, Los Angeles, CA, 139–146.

Rubin, D. B. (1993) Satisfying Confidentiality Constraints Through the Use of Synthetic Multiply Imputed Microdata. *Journal of Official Statistics*, 9, 461–468.

Ting, D., Fienberg, S. E., and Trottini, M. (2005) Unpublished Technical Report.

Trottini, M. (2004) *Decision Models for Data Disclosure Limitation*, Ph.D. thesis, Carnegie Mellon University.

Trottini, M. (2005) *Statistical Disclosure Limitation in Longitudinal Linked Data: Objectives and Attributes*, UNECE/Eurostat Work Session on Statistical Data Confidentiality, Geneva 9-11 November 2005.