Working Paper No. 19 (Summary)

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 ECE/Eurostat work session on statistical data confidentiality
(Luxembourg, 7-9 April 2003)

Topic (v): Risk assessment

ASSESSING DISCLOSURE RISK AND DATA UTILITY: A MULTIPLE OBJECTIVES DECISION PROBLEM

Invited paper

Submitted by the University of Alicante (Spain)[1]

[1] Prepared by Mario Trottini (Mario.trottini@ua.es).

Assessing Disclosure Risk and Data Utility: A Multiple Objectives Decision Problem

Mario Trottini

Dpto. de Estadística e I.O. , Universidad de Alicante, Spain
mario.trottini@ua.es

ABSTRACT

Statistical agencies and Information Organizations that systematically publish statistical data are increasingly concerned with possible misuses of their data release that might lead to disclosure of confidential information about individual respondents represented in the data. Beside legal and ethical considerations, consequences of such disclosure can seriously compromise the trust of the public opinion and thus quality and availability of statistical data, the core of any statistical agency work. As a result of this concern, different transformations of the original data (*data masking)* have been designed aimed to produce new forms of data release that, ideally, should be both *safe*, in the sense prevent disclosure of confidential information, and *valid* in the sense that inferences of legitimate users based on the transformed data agree with inferences that could be made using the original data.

While the number of *data masking* techniques has increased significantly in the last few years, their efficacy has undergone very little exploration and there is an increasing need to design suitable criteria that allow their comparison and identify which one performs best and when. The usual solutions, discussed in the literature of statistical confidentiality and that find many applications in current practice of statistical agencies, consist of two steps:

- Step 1: define suitable measures, S and V that can express the extent to which *safety* and *validity* would be achieved if the *masked data* were released;
- Step 2: define a criterion to compare alternative releases based on the values of S and V.

In the conventional terminology we refer to S and V as *measures of disclosure risk*, and *measures of data utility* respectively and to the criteria in step 2 as *optimality criteria*.
Implementation of the above procedures presents two main difficulties:

(i) Although there is general agreement that optimality criteria should be define on the base of safety and validity associated with the released data, *safety* and *validity* are ambiguous concepts that can not be identified with specific targets. They are multivariate in nature.

For a given disclosure limitation problem, in fact, we can imagine several potential intruders, for each intruder several targets and for a given intruder and a given target several alternative intruder's attacks.

(ii) Even assuming that we have defined suitable measures S and V for the safety and validity, still S and V are usually expressed in different units and have very different meaning so that it is not trivial compare arbitrary pairs (s,v), and (s',v').

The problem is even more complex due to the fact that S and V, from the statistical agency perspective, are random variables. In fact the value that S and V take when a data mask is released, depend on the users' actions that are only partially known to the statistical agency that for example has uncertainty about users prior information, the estimation procedures that they use, etc. Thus each data release induces a distribution over the space of consequences and choosing among alternative release is equivalent to choose among alternative lotteries for (S, V), a much more difficult task than just express preferences over pairs (s, v). Note, for example, that even if S and V are discrete and take few possible values, still the set of distributions for (S, V) is infinite. Existing optimality criteria are essentially of two types:

- Type I: maximize validity under a constraint of minimum safety, i.e. choose the form of data release that maximize validity among those that have safety above a fixed threshold t;

- Type II: build an index or *score* based on S and V.

The use of a threshold value for safety, like in criteria of the type I (the most common case) avoid the problem of different scale between *S* and *V*. However it leads to an "extreme criteria". Suppose for example that both *S* and *V* takes value in *(0,100)* and that the threshold for safety is $t=9$. Suppose you have two alternative releases, R_1 that produces with certainty *(S=9.1, V=90)*, and R_2 that produces with certainty *(S=100,V=89.999999)*. Based on criteria of type I we should select R_1 although many statistical agencies in this situation probably would like to choose R_2 since it is hard to believe that an arbitrary large increase in safety is not worth an infinitesimal decrease in validity. Criteria of type II could be appropriate but should be based on solid theoretical ground. Current implementations are primarily heuristic. This can lead to incoherent criteria. Note that in both cases (criteria of type I and II) the random nature of *S* and *V* is completely neglected thus underestimating the actual agency's uncertainty.

In the terminology used in *multiple objectives decision theory*, step 1 corresponds to define the appropriate set of *attributes* for the decision problem, while step 2 corresponds to assess a suitable *multiattribute utility function*. The best action (form of data release) is the one that maximize the expected utility. Despite the strong connection between the two problems no tentative has been made to use results and techniques from multiple objectives decision theory to implement disclosure limitation problem. This paper explores this possibility. What is proposed is a methodology rather than a universal solution. Research in statistical confidentiality, so far, has been primarily heuristic and this has made difficult quality assessment and comparison of existing procedures. The hope is that the framework presented here can be a first step in that direction. The goal is to provide statistical agencies and researchers in the field with a rational framework that can be used to work together on the problem, discuss and compare ideas using a common language and in a constructive way.