

CHAPTER IV – Methods for Tabular Data

Chapter II presented examples of disclosure limitation techniques used to protect tables and microdata. Chapter III described agency practices in disclosure limitation. This chapter presents more detail concerning methodological issues regarding confidentiality protection in tables.

As mentioned earlier, tables are classified into two categories for purposes of disclosure risk analysis: tables of frequency (or count) data and tables of magnitude data. Tables containing frequency data show the percent of the population that have certain characteristics, or equivalently, the number in the population which have certain characteristics. If a cell has only a few respondents and the characteristics are sufficiently distinctive, then it may be possible for a knowledgeable user to identify the individuals in the population. For tables of frequency data disclosure limitation methods are applied to cells with fewer than a specified **threshold number** of respondents to minimize the risk that individuals can be identified from their data. Disclosure limitation methods applied after tabulation include random rounding, controlled rounding, cell suppression, and controlled tabular adjustment. Disclosure limitation methods applied before tabulation include microdata protection techniques such as data perturbation and data swapping.

Tables of magnitude data typically present the results of surveys of organizations or establishments, where the items published are aggregates of nonnegative reported values. For such surveys the values reported by respondents may vary widely, with some extremely large values and some small values. The confidentiality problem relates to assuring that a person cannot use the published total and other publicly available data to estimate an individual respondent's value too closely. Disclosure limitation methods are applied to cells for which a **linear sensitivity measure** indicates that some respondent's data may be estimated too closely. For tables of magnitude data cell suppression is the most widely used method. Controlled tabular adjustment offers another alternative. Both methods are applied after tabulation. Disclosure limitation methods applied before tabulation include microdata protection techniques such as adding noise.

Tables of frequency data are discussed in Section A. The major methodological areas of interest are in controlled rounding and the use of microdata methods such as data swapping. Tables of magnitude data are discussed in Section B. This section provides some detail concerning linear sensitivity measures, auditing of proposed suppression patterns and automated cell suppression methodologies.

A. Tables of Frequency Data

Tables of frequency data may relate to people or establishments. Frequency data for establishments are generally not considered sensitive because so much information about an establishment is publicly available. Disclosure limitation techniques are generally applied to tables of frequencies based on demographic data. As discussed earlier, the most commonly used **primary disclosure rule** for deciding whether a cell in a table of frequency data reveals too

much information is the "threshold rule". A cell is defined to be sensitive when the number of respondents is less than some predetermined threshold. If there are cells that are identified as being sensitive, steps must be taken to protect them. The methods of preventing disclosure in tables of counts or frequencies were illustrated in II.C.2. Included are: combining cells, random rounding, controlled rounding, cell suppression, controlled tabular adjustment, and microdata techniques. Combining cells is generally a judgmental activity, performed by the survey manager. There are no methodological issues to discuss. Selection of cells for complementary suppression is the same problem for both tables of frequencies and tables of magnitude data. Complementary suppression will be discussed in Section B.2 of this Chapter. Controlled tabular adjustment is most valuable for establishment level data and is also discussed in Section B.2. Microdata techniques have been used to publish data from the decennial census since 1990. These techniques were illustrated in Chapter II, and the technical issues are described in Chapter V.

Perturbation methods include random rounding and controlled rounding as special cases. Controlled rounding is a special case of random rounding. Controlled rounding is the most desirable of the perturbation methods, because it sets a condition that the cell values must add to the published row and column totals. It results in an additive table (sums of row, column and layer entries are equal to the published marginal total). Controlled Rounding can always be solved for two-dimensional tables, and can generally be solved for three-dimensional tables. Section A.1 provides more detail on the methodology used in controlled rounding.

A.1. Controlled Rounding

Controlled rounding was developed to overcome the shortcomings of conventional and random rounding and to combine their desirable features. Examples of random rounding and controlled rounding were given in II.C.2. Like random rounding, controlled rounding replaces an original two-dimensional table by an array whose entries are rounded values that are adjacent to the corresponding original values. However, the rounded array is guaranteed to be additive and can be chosen to minimize any of a class of standard measures of deviation between the original and the rounded tables.

A solution to the controlled rounding problem in two-dimensional tables was found in the early 1980's (Cox and Ernst, 1982). With this solution the table structure is described as a mathematical network, a linear programming method that takes advantage of the special structures in a system of data tables. The network method can also be used to solve controlled rounding for sets of two-dimensional tables that are related hierarchically along one dimension (Cox and George, 1989).

For three-dimensional tables an exact network solution does not exist (Cox and Ernst, 1982). Current methods make use of an iterative approximate solution using a sequence of two-dimensional networks. The exact solutions for two-dimensional tables and the approximate solutions for three-dimensional tables are both fast and accurate. Although solutions to the controlled rounding problem are available, controlled rounding is not a common practice among

U.S. government agencies.

B. Tables of Magnitude Data

For tables of magnitude data the values reported by respondents are aggregated in the cells of a table. Examples of magnitude data are income for individuals and sales volumes and revenues for establishments. Particularly for establishments, these reported values are typically highly skewed with a few very large reported values that might easily be associated with a particular respondent by a knowledgeable user. As a result, a more mathematical definition of a **sensitive cell** is needed for tables of magnitude data. For tables of frequency data each respondent contributes equally to each cell, leading to the simple threshold definition of a sensitive cell.

Mathematical definitions of sensitive cells are discussed in Section B.1 below. After the tables have been created and the sensitive cells are identified, a decision must be made as to how to prevent disclosure from occurring. For tables of magnitude data the possibilities include combining cells and rolling up categories, cell suppression, and controlled tabular adjustment. All were summarized and illustrated in Chapter II.

In the combination method tables are redesigned (categories rolled-up) so there are fewer sensitive cells. Table redesign methods are useful exercises, particularly with tables from a new survey or where portions of a table contain many sensitive cells because the population is sparse. However, it is not generally possible to eliminate all sensitive cells by collapsing tables, and rigorous automated procedures for collapsing in general remain to be developed.

The historical method of protecting sensitive cells in tables of magnitude data is cell suppression. Sensitive cells are not published (they are suppressed). These sensitive suppressed cells are called **primary suppressions**. To make sure the primary suppressions cannot be derived by subtraction from published marginal totals, additional cells are selected for **complementary suppression**. Complementary suppressions are sometimes called **secondary suppressions**.

For small tables, it is possible to manually select cells for complementary suppression, and to apply audit procedures (see Section 2.a) to guarantee that the selected cells adequately protect the sensitive cells. For large-scale survey publications with many related tables, the selection of a set of complementary suppression cells that are "optimal" in some sense is an extremely complex problem. Complementary suppression is discussed in Section B.2.

Controlled tabular adjustment is also illustrated in Chapter II. Some of the technical details are discussed in IV.B.3. Finally, microdata methods are increasingly being used to protect tabular data prior to tabulation. For establishment level data, noise addition is the technique that has been applied to date. This is summarized in Chapter II, and discussed in more detail in IV.B.4.

B.1. Definition of Sensitive Cells – Linear Sensitivity Rules

The definitions and mathematical properties of linear sensitivity measures and their relationship to the identification of sensitive cells in tables were formalized by Cox (1981). Although the common linear sensitivity rules were known in 1978 and were used to identify sensitive cells, their mathematical properties had not been formally demonstrated. The important definitions and properties are given below.

For a given cell, X , with N respondents the respondent level data contributing to that cell can be arranged in order from large to small: $x_1 \geq x_2 \geq \dots \geq x_N \geq 0$. Then, an **upper linear sensitivity measure**, $S(X)$, is a linear combination

$$S(X) = \sum_{i=1}^N w_i x_i$$

defined for each cell or cell union X and its respondent data $\{x_i\}$. The sequence of constants, $\{w_i\}$, is called the sequence of weights of $S(X)$. These weights may be positive or negative. A cell or cell union X is **sensitive** if $S(X) > 0$. Note that multiplying a linear sensitivity measure by a constant yields another (equivalent) linear sensitivity measure. The linear sensitivity measures described in this section are all normalized so that the weight multiplying x_1 is equal to 1. This normalization makes it easier to compare them. If a respondent contributes to two cells, X and Y , then it remains a single respondent to the union of X and Y , with value equal to the sum of its X and Y contributions.

One of the properties which assists in the search for complementary cells is **subadditivity**, which guarantees that the union of disjoint cells which are not sensitive is also not sensitive. Cox shows that a linear sensitivity measure is subadditive if the sequence of weights is nonincreasing, i.e. if $w_1 \geq w_2 \geq \dots \geq w_N$. Subadditivity is an important property because it means that aggregates of cells which are not sensitive are not sensitive and do not need to be tested. Valid complementary cells have the property that their union with the sensitive cell(s) in a row, column or layer where marginal totals are published is not sensitive according to the linear sensitivity measure. A simple result is that zero cells are not valid candidates for complementary suppression as the union of a sensitive cell and a zero cell is equal to the sensitive cell, and is therefore still sensitive. Complementary suppressions may not be needed if marginal totals are not published.

The commonly used primary suppression rules are described Sections a, b, and c below. They are compared in Section d. Each of these rules involves parameters that determine the values taken by the weights, $w_1 \dots w_n$. Although agencies may reveal the primary suppression rule they use, they should not disclose parameter values, as knowledge of the rule and its parameters enables a respondent to make better inferences concerning the values reported by other respondents. An example is presented in Section 3.

There are three linear sensitivity measures that are discussed in the literature and used in practical applications. These are the p-percent rule, the pq rule and the (n, k) rule. They are described below. All are subadditive, as can be seen by the fact that the weights in the equations defining $S(x)$ are non-increasing. The p-percent and pq rule classify cells of count data as sensitive if $n < 3$.

B.1.a. The p-Percent Rule

Approximate disclosure of magnitude data occurs if the user can estimate the reported value of some respondent too accurately. Such disclosure occurs, and the table cell is declared sensitive, if upper and lower estimates for the respondent's value are closer to the reported value than a pre-specified percentage, p . This is referred to as the "p-percent estimation equivocation level" in Statistical Policy Working Paper 2. It is more generally referred to as the **p-percent rule**, and has linear sensitivity measure,

$$S^{p\%}(X) = x_1 - \frac{100}{p} \sum_{i=c+2}^N x_i.$$

Here, c is the size of a coalition, a group of respondents who pool their data in an attempt to estimate the largest reported value. The cell is sensitive if $S^{p\%}(X) > 0$. Note that if there are less than 3 respondents ($N < 3$) in cell X , then $S^{p\%}(X) = x_1 > 0$ and the cell is sensitive for all values of p and c .

The p-percent rule is derived as follows. Assume that from general knowledge any respondent can estimate the contribution of another respondent to within 100-percent of its value. This means that the estimating respondent knows that the other respondents' values are nonnegative and less than two times the actual value. For the p-percent rule, it is desired that after the data are published no respondent's value should be estimable more accurately than within p percent (where $p < 100$).

It can be shown that the coalition including the second largest respondent is in a position to estimate the value of x_1 most accurately, and that if x_1 is protected, so are all the smaller respondents. Thus, it suffices to provide the protection to the largest respondent, and to assume that the estimating party is a coalition of the second largest respondent and the next largest $c - 1$ respondents. As the coalition respondents may estimate each of x_{c+2}, \dots, x_N to within 100 percent, they have an estimate for the sum of these smallest respondents, E , such that

$$\left| \sum_{i=c+2}^N x_i - E \right| \leq \sum_{i=c+2}^N x_i.$$

They can estimate the value of x_1 by subtracting the value they reported to the survey $\left(\sum_{i=2}^{c+1} x_i\right)$ and their estimate of the smaller respondent's total, E , from the published total. The error in this estimate will be equal to the error in estimating E , which is less than or equal to $\sum_{i=c+2}^N x_i$. The requirement that this estimate be no closer than p -percent of the value of x_1 ($p < 100$) implies that

$$\sum_{i=c+2}^N x_i \geq \frac{p}{100} x_1.$$

This can be rewritten as the linear sensitivity rule above. A simpler version of the p percent rule that assumes coalitions of size c can be written as follows:

$$S = x_1 - 100/p * (T - T_c - x_1)$$

Where T is the cell total of all respondents, T_c is the sum of the respondent values in the coalition, and x_1 is the largest value. Using this formula the cell is sensitive if S is positive. In the simple case where $T_c = x_2$ (i.e., the coalition is only a size of one), then $T - T_c - x_1 = T - x_2 - x_1$ which means the remaining cell value is the sum of all the smallest companies in the cell with the exception of the two largest. $T - T_c - x_1$ will equal zero only if the coalition (T_c) includes all the respondents in the cell except the largest company.

B.1.b. The pq Rule

In the derivation for the p -percent rule, we assumed that there was limited prior knowledge about respondent's values. Some people believe that agencies should not make this assumption. In the pq rule, agencies can specify how much prior knowledge there is by assigning a value q which represents how accurately respondents can estimate another respondent's value before any data are published ($p < q < 100$). Thus, there is an improved estimate, E_2 , of $\sum_{i=c+2}^N x_i$ with the property that

$$\left| \sum_{i=c+2}^N x_i - E_2 \right| \leq \frac{q}{100} \sum_{i=c+2}^N x_i.$$

This leads directly to a more accurate estimate for the largest respondent's value, x_1 . The requirement that this estimate be no closer than p -percent of the value of x_1 implies that

$$\frac{q}{100} \sum_{i=c+2}^N x_i \geq \frac{p}{100} x_1.$$

$$S^{pq}(X) = x_1 - \frac{q}{p} \sum_{i=c+2}^N x_i.$$

This can be rewritten as the linear sensitivity rule.

Note that the pq rule (sometimes called a prior-posterior ambiguity rule) and the p-percent rule are identical if the ratio q/p , the "information gain", is equal to $100/p$. In the table below we use the ratio q/p as a single parameter for the pq rule. If users fix a value for p and a value for $q < 100$, the pq rule is more conservative (will suppress more cells) than the p-percent rule using the same value of p .

Note that if there are fewer than 3 respondents ($N < 3$), then $S^{pq} = x_1 > 0$ and cell X is sensitive for all values of c and q/p .

Most frequently the pq rule is given with the size of a coalition equal to 1. In this case the linear sensitivity rule is given by

$$S^{pq}(X) = x_1 - \frac{q}{p} \sum_{i=3}^N x_i.$$

B.1.c. The (n, k) Rule

The **(n, k) rule**, or dominance rule was described as follows in Statistical Policy Working Paper 2. "Regardless of the number of respondents in a cell, if a small number (n or fewer) of these respondents contribute a large percentage (k percent or more) of the total cell value, then the so-called **n respondent, k percent rule** of cell dominance defines this cell as sensitive." Many people consider this to be an intuitively appealing rule, because, for example, if a cell is dominated by one respondent then the published total alone is a natural upper estimate for the largest respondent's value. Although coalitions are not specifically discussed in the derivation of the (n, k) rule, agencies select the value of n to be larger than the number of any suspected coalitions. Many agencies use an (n, k) rule with $n = 1$ or 2.

The linear sensitivity measure for the (n, k) rule is given by

$$S^{(n,k)}(X) = \sum_{i=1}^n x_i - \frac{k}{100-k} \sum_{i=n+1}^N x_i.$$

If $N \leq n$, $S^{(n,k)} = \sum_{i=1}^N x_i > 0$ and cell X is sensitive for all values of k .

B.1.d. The Relationship Between (n, k) and p-Percent or pq Rules

Table 1 is designed to assist users in selecting a value of the parameter p for use with the p -percent rule with coalitions of size 1 (or for the value of the ratio, q/p , for the pq rule with coalitions of size 1) when they are used to thinking in terms of the (n, k) rule. For various values of $p\%$ (q/p), the table shows the value of k_1 and k_2 such that if the linear sensitivity rule for $(1, k_1)$ or $(2, k_2)$ is positive then the linear sensitivity rule for the p -percent (pq) rule will be positive. With this formulation, the p -percent (pq) rule is more conservative. It will suppress more cells than will either of the two (n, k) rules individually, and also more than the combination rule based on the two (n, k) rules. The derivation of the inequalities used in Table 1 is presented in the Technical Notes at the end of this Chapter. Additionally, the sensitivity regions for (n, k) , p -percent, and pq rules are illustrated graphically in the Technical Notes. See Robertson, D. A. (1993) for theoretical analysis comparing disclosure rules.

To illustrate the use of Table 1, if the analyst wants to make sure that a cell where the largest respondent contributes more than 75 percent of the total is suppressed, and that a cell where the largest two respondents exceed 85 percent of the total is suppressed, he/she could approximately accomplish this by using the p -percent rule with p equal to 33.3 percent, or the pq rule with information gain, $q/p = 3$.

The p -percent, pq and (n, k) rules as well as the combination rule,

$$S^{comb} = \max(S^a(X), S^b(X))$$

are subadditive linear sensitivity rules. (Here $S^a(X)$ and $S^b(X)$ denote any two subadditive linear sensitivity measures.) Any of these rules is acceptable from a mathematical point of view. However, the p -percent or pq rule is preferred for two major reasons. First, the tolerance interval concept directly parallels methods currently used for complementary suppression, (see section B.2.a.iii). Second, as illustrated in the table above and the example in the Technical Notes, the p -percent (pq) rule provides more consistent protection areas than a single version of the (n, k) rule.

TABLE 1 Relationship Between Suppression Regions for p-Percent or (pq) Rule and (1,k), (2,k) Rules

		$S^{p\%}(X) > 0$ and	Sensitive when cell
P	q/p	x_1/T exceeds:	$(x_1 + x_2)/T$ exceeds:
50.0%	2	66.7%	80.0%
33.3%	3	75.0%	85.7%
16.7%	6	85.7%	92.3%
11.1%	9	90.0%	94.7%

NOTE: $T = \sum_{i=1}^N x_i$ is the cell total.

B.1.e. Information in Parameter Values

Agencies may publish their suppression rules, however, they should keep the parameter values they use confidential. Knowledge of both the rule and the parameter values enables the user to make better inferences concerning the value of suppressed cells, and may defeat the purpose of suppression.

For example, assume that an agency uses the p-percent rule with p=20 percent, and that the same value of p is used to determine the protection regions for complementary suppression. We assume that a cell total is 100 and that the cell is sensitive according to the p-percent rule. That cell will be suppressed along with other suitable complementary cells. For this cell (as with any suppressed cell), any user can use a linear programming package to calculate upper and lower bounds for the cell total based on the published row and column equations. Assume that this leads to the following inequality:

$$80 = \text{lower bound} \leq \text{cell total} \leq \text{upper bound} = 120.$$

In this case, the protection region used in selecting cells for complementary suppression assures that the cell total cannot be estimated more closely than plus or minus 20 percent of the cell value, or plus or minus 20 in this case. A knowledgeable user has thus uniquely determined that the value of the suppressed cell total must be 100. Once the total for one suppressed cell has been uniquely determined, it is likely that other cell values can easily be derived by subtraction from published marginal totals.

B.2. Complementary Suppression

Once sensitive cells are identified by a primary suppression rule, other nonsensitive cells must be selected for suppression to assure that the respondent level data in sensitive cells cannot be estimated too accurately. This is the only requirement for a proposed set of complementary cells for tables of magnitude data and is generally considered to mean that a respondent's data cannot be estimated more closely than plus or minus some percentage.

There are two ways respondent level data can be compromised. First, implicitly published unions of suppressed cells may be sensitive according to the linear sensitivity measure. This depends on the characteristics of the respondent level data in the cell union, and tends to be a problem only where the same respondents contribute to both cells. Second, the row and column equations represented by the published table may be solved, and the value for a suppressed cell estimated too accurately. Automated methods of **auditing** a proposed suppression pattern may be needed to assure that the primary suppressions are sufficiently well protected (see Section B.2.a).

Any set of cells proposed for complementary suppression is acceptable as long as the sensitive cells are protected. For small tables this means that selection of complementary cells may be done manually. Typically the data analyst knows which cells are of greatest interest to users (and should not be used for complementary suppression if possible), and which are of less interest to users (and therefore likely candidates for complementary suppression). Manual selection of complementary cells is acceptable as long as the resultant table provides sufficient protection to the sensitive cells. An automated audit should be applied to assure this is true.

For large systems of tables, for example, those based on an Economic Census, the selection of complementary cells is a major effort. Manual selection of cells may mean that a sensitive cell is inadvertently left unprotected or that consistency is not achieved from one table to another in a publication. (Cox, 1980). Inconsistency in the suppression patterns in a publication increases the likelihood of inadvertent disclosure. For this reason linear programming techniques have been applied to the selection of cells for complementary suppression by statistical agencies for many years. (Cox, 1995). As an additional benefit, agencies expect automated selection of the complementary cells will result in less information lost through suppression. Examples of the theory and methods for automatic selection of cells for complementary suppression are discussed in Section B.2.b.

B.2.a. Audits of Proposed Complementary Suppressions

B.2.a.i. Implicitly Published Unions of Suppressed Cells Are Sensitive

If sensitive cells are protected by suppressing other internal table cells when publishing the marginal totals, the implicit result is that the unions of the suppressed cells in rows, columns and layers are revealed by subtracting from the total. Thus, one way to audit the protection supplied by the suppression pattern is to apply the linear sensitivity rule to those unions to assure that they are not sensitive. While this type of audit is a simple matter for small tables, Cox (1980) points

out that for large tables it may be computationally intractable unless a systematic approach is used. This type of audit is not included in standard audit software because of its dependence on respondent level data.

Clearly a table for which suppression patterns have been selected manually requires an audit to assure that the pattern is acceptable. Early versions of complementary suppression software used approximation arguments to select cells for complementary suppression (individual respondent data were not used.) These methods guaranteed that unions of suppressed cells were not sensitive as long as different respondents contributed to each cell. However, if the same respondents contributed to multiple cells in a cell union, then an audit was needed.

B.2.a.ii. Row, Column and/or Layer Equations Can Be Solved for Suppressed Cells

A two-dimensional table with row and column subtotals and a three-dimensional table with row, column and layer subtotals can be viewed as a large system of linear equations. The suppressed cells represent unknown values in the equations. It is possible that the equations can be manipulated and the suppressed values estimated too accurately. Audits for this type of disclosure require the use of linear programming techniques. The output of this type of audit is the maximum and the minimum value each suppressed cell can take given the other information in the table. When the maximum and the minimum are equal, the value of the cell is disclosed exactly. To assure that cells cannot be estimated too accurately the analyst makes sure the maximum and the minimum value for the suppressed cell are no closer to the true value than some specified percentage protection.

It is well known that a minimal suppression pattern where marginal totals are presented will have at least two suppressed cells in every row, column and layer requiring suppression. This is not sufficient, however, as was illustrated in Chapter 2 Section C.2.a.

B.2.a.iii. Software For Auditing A Suppression Pattern

Automated methods of auditing a suppression pattern have been available since the mid 1970's at the U.S. Census Bureau, and at Statistics Canada. Modern versions of audit software set up the linear programming problem and use commercially available linear programming packages. All audit systems produce upper and lower estimates for the value of each suppressed cell based on linear combinations of the published cells. A suppression audit can uncover three types of problems for tables cells: 1) the upper and lower limits may be the same; 2) the upper and lower limits may be too close together; 3) the upper and/or lower limits may be too close to the cell value. The data analyst uses the output from the audit to determine whether the protection provided to the sensitive cells by the proposed complementary cells is sufficient. The user should know the type and extent of the rounding of cell values in a table that is being audited to avoid misleading evaluations of data protection. Depending upon whether suppression was

applied to the rounded or unrounded data can result in over- or under-suppression of cells in a table. These audit methods are applicable to tables of both magnitude and frequency.

Linear programming is the most common procedure used for auditing suppression patterns in a table because it can be used for higher-dimensional tables. (Zayatz 1992a). Network procedures have been shown to provide fast solutions for two-dimensional tables. The network flow method for cell suppression is self-auditing only for two-dimensional tables in which there is a hierarchy in one dimension. The network flow method is not-self auditing for two-dimensional tables with a hierarchical variable structure in both the row and column, and it is not self-auditing for three dimensional or higher dimensional tables that contain a hierarchical structure. (Massell 2002).

At the U. S. Census Bureau both types of audits are subsumed into the algorithm that selects cells for complementary suppression. The company level contributions for a cell are used in selecting a protection level or tolerance interval for each cell that provides protection to all respondents in the cell. The algorithm that selects cells for complementary suppression provides that the primary cells cannot be estimated more accurately than that specified tolerance interval. The complementary suppressions selected by applying the algorithm do not require additional audits.

Audit software was developed by the Confidentiality and Data Access Committee, with support from a number of statistical agencies, and is available with documentation at <http://www.fcs.m.gov/committees/cdac/resources.html>. This software is written in SAS[®] and checks the lower and upper bounds around suppressed cells in a table that contains non-additive, independently rounded cells. The program requires a specific format for the ASCII input file. The program also checks for whether independent rounded cells exist in the table and adjusts the cell values to preserve additivity within the row and columns at the same time it is performing the import function. The user has the option of specifying a protection range based on a plus/minus percent basis or absolute value basis. The software is not limited by the number of dimensions in a table and the linear programming methodology provides for two types of optimizers.

B.2.b. Automatic Selection of Cells for Complementary Suppression

Software that automatically selects complementary cells for suppression has been available since the 1970's at Statistics Canada and at the U.S. Census Bureau. These programs typically use linear programming methods implemented by accessing general purpose linear programming routines which make use of special structures in the data. Due to refinements in linear programming algorithms, these routines run much faster now than in the 1980's. Network flow methods may be viewed as a special case of linear programming. They work best for two dimensional tables, with at most one level of hierarchy (in either rows or columns). Routines based on network flow methods are typically much faster than linear programming routines. Cell suppression programs can be used for both magnitude data tables and frequency data tables.

At the U.S. Census Bureau these programs have been used mainly for business survey magnitude data. Robert Jewett (Jewett, 1993) wrote a set of cell suppression programs for this purpose. They include much beyond the basic cell suppression model. For example, the program can be used to identify sensitive cells from a given input microdata file by using the $p\%$ rule. The problem of “common respondents” is handled by defining a table of capacities for each primary. It is constructed just before complementary suppressions are selected to protect a given primary. The “common respondents” problem arises frequently with business survey data since many companies have more than one establishment and often these establishments are contributors to different cells of the same table. The U.S. Census Bureau must protect not only each establishment’s contribution but all sums of an establishment’s facilities, including the company’s total contribution. These programs are also able to handle tables that are linked and interrelated to cells in two or more of the tables. It uses the method of backtracking to check that a given suppressed cell has the same degree of uncertainty in each table in which it appears.

The software, tau-Argus, developed from the Computational Aspects of Statistical Confidentiality (CASC) European project offers methods to identify sensitive cells, a choice of algorithms to select secondary suppressions, an suppression audit program to compute interval bounds for suppressed cells, and a module to generate synthetic values to replace suppressed original ones in a publication. The documentation and software for operating tau-Argus are available at <http://neon.vb.cbs.nl/casc>.

In the straightforward implementation of linear programming, sensitive cells are treated sequentially beginning with the most sensitive. At each step (i.e. for each sensitive cell) the set of complementary cells that minimizes a cost function (usually the sum of the suppressed values) is identified. Zayatz (1992a) describes the formulation for two-dimensional tables. Zayatz (1992b) gives the parallel formulation for three-dimensional tables. As above, these are implemented by using a commercially available linear programming package. The disadvantage of the straightforward linear programming approach is the computer time it requires. For large problems, the run time of the Central Processing Unit of a personal computer increases significantly with 3 or more dimensions.

Another linear programming approach is based on describing the table structure as a mathematical network, and using that framework and the required tolerance intervals for each cell to balance the table. The network methods are favored because they give the same result as the straightforward linear programming methods, but the solution requires much less computer time. The network method is directly applicable to two-dimensional tables and to two-dimensional tables with subtotal constraints in one dimension (Cox, 1995). Subtotal constraints occur when data in one dimension have a hierarchical additive structure such as the North American Industry Classification System (NAICS) coding system. In the past 20 years, there was considerable research in developing faster and more efficient procedures for both two-dimensional and three-dimensional tables. Research has involved using methods based on integer programming, network flow theory, and neural networks.

Complementary suppression and controlled rounding can both be solved using network theory. Ernst (1989) demonstrated the impossibility of representing a general three or higher dimension table as a network. For this reason, complementary suppression for three-dimensional tables currently uses linear programming as the main approach. (Zayatz, 1992b). The straightforward linear programming methods can be used for small three-dimensional tables. However, for large three-dimensional tables, an iterative approximate approach based on a sequence of two-dimensional networks is used. The complementary suppression pattern identified by this approximate approach must still be audited to assure that an individual sensitive cell cannot be estimated too accurately.

As mentioned above, one possible objective or cost function for automated procedures is to minimize the sum of the suppressed values. With this objective function, automated procedures tend to suppress many small cells, a result not generally considered "optimal" by the analyst. Other possible cost functions include minimizing the total number of suppressed cells in a table or minimizing the suppression for specific data series in a table. Further research is needed into the identification of cost functions for use in selecting the "optimal" complementary suppressions. Possibilities here include research into a cost function for use in a single run of the software, as well as cost functions for use in multiple runs of the software. An example is the development of a cost function that is used during a second pass through the software to remove superfluous suppressions (Zayatz, 1992b).

Another reason the complementary cells selected by automated methods do not provide the "optimal" set for the table as a whole is that all current implementations protect sensitive cells sequentially. For any given sensitive cell, the complementary cells selected to protect it will be optimal according to the objective function, conditional on all suppressions selected for previously considered sensitive cells. The sequential nature of the approach leads to over-suppression.

In spite of the lack of "optimality" of the result, the automated complementary cell suppression procedures identify useful sets of complementary suppressions. However, work is often needed to fine tune, reduce over-suppression, and assure that the analysts' nonmathematical definition of an "optimal" solution is more closely realized.

B.3. Controlled Tabular Adjustment

Controlled Tabular Adjustment is a useful methodology for protecting tables of magnitude data as well as count data. It is discussed with an example in Chapter 2 Section D.3.d. Each sensitive original value in a table is replaced with an imputed safe value that is a sufficient distance from the true sensitive value. Some of the remaining non-sensitive cell values are adjusted from their true values by as small an amount as possible to restore additivity to the published totals. CTA can be applied to produce solutions where marginal sums are minimally changed. However, allowing minor adjustments to the marginal values reduces the need for larger adjustments to the internal non-sensitive cells in a table.

There are two different approaches that apply CTA methodology. The original CTA method uses a linear programming method to restore additivity to the table. Initially, the **LP-based Controlled Tabular Adjustment** procedure used the reciprocal of the cell values as a cost function to minimize the overall deviation of non-sensitive cells from the true cell value. Another appropriate optimization function may be to minimize the sum of the absolute values of the data adjustments. The reciprocal of the cell value allows for larger changes to large cells and causes smaller changes to small cells when compared with other cost functions. Most LP based procedures review the solution quality and feasibility using the underlying table structure. The algorithm systematically changes sensitive and nonsensitive cells first seeking to obtain a feasible solution, and then once feasibility is reached, then it moves on to optimize the quality of the adjustment using a pre-specified cost function. Software that use some type of adaptive memory process for reviewing the optimal adjustments provide better results in terms minimal adjustments to cell values than those methods that apply a “rigid memory” design such as a branch and bound technique.

During the first phase of applying either type of CTA methodology, the sensitive cells are ordered from largest to the smallest. By using an alternating sequence, the ordered sensitive cell values are then changed to either lower or upper protection bounds. After completing the changes to all the sensitive cells in the table, non-sensitive table cells are considered to restore the additive table structure.

A second approach, called **simplified Controlled Tabular Adjustment**, was developed as a cost effective alternative to the original LP-based CTA method. The simplified CTA minimizes the percentage deviation from the true cell value for non-sensitive cells as its optimization function. The minimum percent deviation criteria used in simplified controlled tabular adjustment produces similar results as the reciprocal of the cell value-based cost function used in the LP-based approach. (Dandekar, 2004). Simplified CTA is easier to implement and more computationally efficient than the LP-based CTA procedure, although further research is needed on different table structures to further evaluate these two approaches. LP based CTA and simplified CTA use different approaches to restore additivity to the table structure. The original CTA method uses a linear programming method to restore table additivity. The simplified CTA method, on the other hand, accepts all necessary adjustments in marginal table cell values to restore additive table structure.

B.4. Adding Noise to Microdata Prior to Tabulating Data

Adding noise to the underlying microdata is a method that has been used to protect magnitude tabular data. It is different from the noise procedures used to protect public use microdata files. The noise addition method adjusts each value by a small amount (the exact percent to remain confidential within the statistical agency). Each establishment reporting in the sample or survey is assigned a multiplier, or noise factor. A company may have several different stores or establishments. In this case, each establishment may be assigned a slightly different multiplier as

long as the overall distribution of the multipliers across all establishments within a company average the specified percent for adjusting that company's reported values. (Evans, 1998).

For example, if an establishment's data is adjusted by 10%, then its data would be multiplied by a number that is close to either 1.1 or 0.9. Any type of distribution can be used to choose the multipliers for each establishment. In this example, whatever distribution is used to generate a multiplier of 1.1, it is important that the same distribution shape, or its "mirror image," be used to generate the multipliers near 0.9 to adjust data in the opposite direction. The two distributions of multipliers should produce a joint distribution of multipliers that is symmetrical and approximates 1.

The direction of adding the noise to each responding company is randomly assigned. Using the example of 10% as the base for perturbation, this is equivalent to determining if all establishments in a company have multipliers close to 1.1 or close to 0.9. The next step in the process is to randomly assign a multiplier to each establishment within a company. The multipliers would be generated from that half of the overall distribution of the multipliers that corresponds to the direction of perturbation assigned to that company. An example of assigning multipliers to a set of respondents is as follows:

Example 1:

<u>Company</u>	<u>Establishment</u>	<u>Direction</u>	<u>Multiplier</u>
Company A		1.1	
	Establishment A1		1.12
	Establishment A2		1.09
	Establishment A3		1.10
	Establishment A4		1.11
Company B		0.9	
	Establishment B1		0.89
	Establishment B2		0.93
Company C		1.1	
	Establishment C1		1.08

In this example, the expected value of the amount of noise added in any cell value is zero because of the symmetry of the distribution of the multipliers and the random assignment of both the direction of perturbation and the multipliers within each company. The probability that a company's establishments will be perturbed in a positive direction is equal to the probability that they will be perturbed in a negative direction. The distribution of the multipliers is symmetric about 1. The expected value of any given multiplier is 1, hence the expected value of the *amount* of noise in any given establishment is 0, and the amount of noise in any cell value is simply the sum of the noise in its component establishments.

Noise addition differs from Controlled Tabular Adjustment because noise addition adjusts the reported values prior to any tabulations. Controlled Tabular Adjustment adjusts the cells after

the data have been tabulated on a cell by cell basis. Noise addition relies on the random assignment of the multiplier to control the effects of adding noise to different types of cells.

C. Online Data Query Systems

Most online query systems that were developed by the federal agencies allow access to summary files with matrices of aggregated data. These query systems allow users to design queries to generate customized tabulations. Special disclosure limitation methods should be considered when users access microdata files to produce customized tabulations.

One example of an online query system that allows users to access microdata files is the “Advanced Query System” (AQS) which is part of the Census Bureau’s “American Fact Finder” online data dissemination system. The microdata files in the AQS contain information on individuals and households. To ensure that tabulations from these microdata files do not reveal the identities of respondents, the Census Bureau uses data recoding and data swapping techniques in addition to other microdata techniques.

Variables such as geography, detailed race, age, occupation, industry, Hispanic origin, and group quarters are re-coded and/or collapsed. All continuous variables, such as income, fuel and utility costs, property taxes, rent, and mortgage payments are top coded to mask the outlying values in the tails of the distributions of each continuous variable. The re-coded variables are added to the files used by AQS. An external user is diverted to the re-coded variables and geographic area when submitting a query.

In addition to recoding, a swapping technique is also applied to the records in the microdata files. The technique consists of swapping pairs of household records selected as having the highest disclosure risk based upon a predetermined set of key variables. In the AQS system, records are selected for swapping with a probability inversely proportional to block size.

Any request submitted by a user passes through two filters; the Query filter and the Statistical Results filter. The purpose of the Query filter is to detect those queries that will not pass disclosure limitation before the query is submitted for execution, such as the geographic variable must meet a minimum threshold. The Statistical Results Filter checks the final values in the cells of the resulting table. If a table does not pass the filters, the entire table is suppressed and the user does not receive the table. The AQS system does not perform any cell suppression. A message is sent to a user that the table is suppressed for confidentiality reasons and the user may then try requesting a table with less detail.

The disclosure protection procedures applied by the Agriculture Resource Management Survey (ARMS) uses a different approach than the AQS system. The ARMS on line query system allows users to select across survey data sets and build customized reports. There are three stages to the disclosure protection procedures used in the ARMS system. In the first step, noise is added to the weights for underlying microdata in a unique way to protect large establishments that may dominate a cell. The second step is to develop minimum expanded farm counts in a

cell and to test the sensitivity of that cell using the p-percent rule. The third step applies primary cell suppression without any complementary suppression. No complementary suppression is necessary because of the noise that was initially added to the microdata provides the necessary protection to the aggregates. Cells in the outputted files are suppressed if they fail any of the three criteria: 1) if the ratio of the cell value with noise to the cell value without noise is outside a set range, then the cell is suppressed; 2) if the weighted farm count for a cell is small then the cell is suppressed; 3) if the cell fails the p-percent rule and has insufficient noise to protect the actual value, then the cell is also suppressed. The approach used in ARMS avoids the need for complementary suppression and simplifies the computational problems associated with disclosure protection in an on-line query system.

D. Technical Notes: Relationships Between Common Linear Sensitivity Measures

This section illustrates the relationship between the p-percent, pq and (n, k) rules described in the text by using plots of regions of cell sensitivity. To simplify this presentation we make a few assumptions. First, for the p-percent rule we assume there are no coalitions ($c = 1$) and for the (n, k) rules we consider only $n = 1$ and $n = 2$. Second, replace $\sum_{i=3}^N x_i$ by $(T - x_1 - x_2)$. Third, divide each sensitivity rule through by the cell total, T , and multiply by 100. Finally, set $z_i = 100x_i / T$, the percent contributed to the cell total by company i . The sensitivity rules can be written

$$S^{p\%}(X) = \left(1 + \frac{100}{p}\right) z_1 + \frac{100}{p} z_2 - \frac{100}{p} 100,$$

$$S^{pq}(X) = \left(1 + \frac{q}{p}\right) z_1 + \frac{q}{p} z_2 - \frac{q}{p} 100,$$

$$S^{(1,k_1)}(X) = \left(1 + \frac{k_1}{100 - k_1}\right) z_1 - \frac{k_1}{100 - k_1} 100$$

$$S^{(2,k_2)} = \left(1 + \frac{k_2}{100 - k_2}\right) z_1 + \left(1 + \frac{k_2}{100 - k_2}\right) z_2 - \frac{k_2}{100 - k_2} 100$$

The regions where these sensitivity rules are positive (i.e. where the cells are sensitive) are shown in Figure 1. The horizontal axis represents the percent contributed by the largest unit, z_1 and the vertical axis represents the percent contributed by the second largest unit, z_2 . Since $z_1 \geq z_2$ and $z_1 + z_2 \leq 1$ (the sum of the two largest is less than or equal to the cell total), the only possible values in a table cell will be in the lower triangular region bounded from below by the line $z_2 = 0$, from above by the line $z_1 = z_2$ and to the right by the line $z_1 + z_2 = 1$.

The $(1, k_1)$ and $(2, k_2)$ rules are particularly simple to illustrate graphically. The inequality $(1, k_1)$ rule simplifies, and a cell is classified as sensitive if $z_1 > k_1$. The dividing line between sensitive and nonsensitive region is given by a vertical line through the point $(0, k_1)$. Similarly, the inequality for the $(2, k_2)$ rule simplifies and a cell is classified as sensitive if $(z_1 + z_2) > k_2$. The dividing line between the sensitive and nonsensitive regions is the line through the points $(0, k_2)$ and $(k_2, 0)$. This line intersects $z_1 = z_2$ at the point $(k_2/2, k_2/2)$. In all cases the sensitive region is the area to the right of the dividing line. The sensitivity regions for the $(1, 75)$ and $(2, 85)$ rules are illustrated in Figure 1A.

For the p-percent rule the inequality above yields the boundary line for sensitive cells as the line joining the points $(0, 100)$ and $\left(\frac{100}{\frac{p}{100} + 1}, 0\right)$. This line intersects $z_1 = z_2$ at the point

$\left(\frac{100}{\frac{p}{100} + 2}, \frac{100}{\frac{p}{100} + 2}\right)$. The pq rule is the same, with $q/p = 100/p$.

FIGURE 1A
 EXAMPLES OF SUPPRESSION REGIONS
 THE (N,K) RULE WITH N=1 AND K=75, N=2 AND K=85

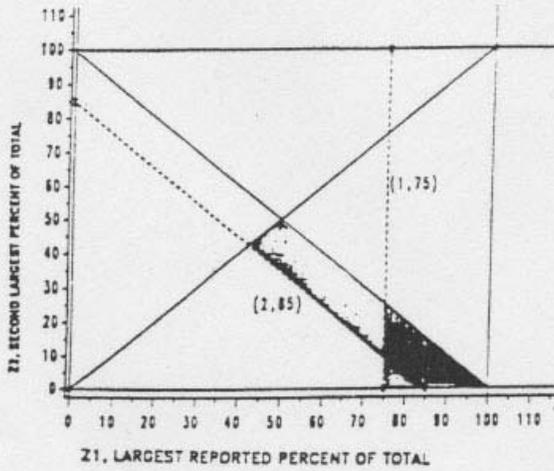


FIGURE 1B
 EXAMPLES OF SUPPRESSION REGIONS
 THE P-PERCENT RULE WITH P=17.65 PERCENT, AND P=35.3 PERCENT

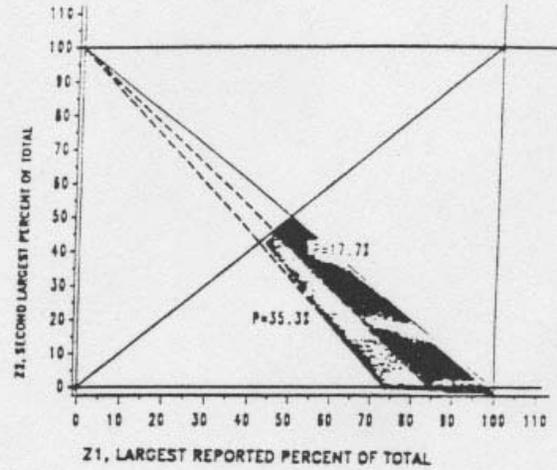


FIGURE 1C
 P-PERCENT LESS CONSERVATIVE THAN (2,85)
 P = 17.7 PERCENT

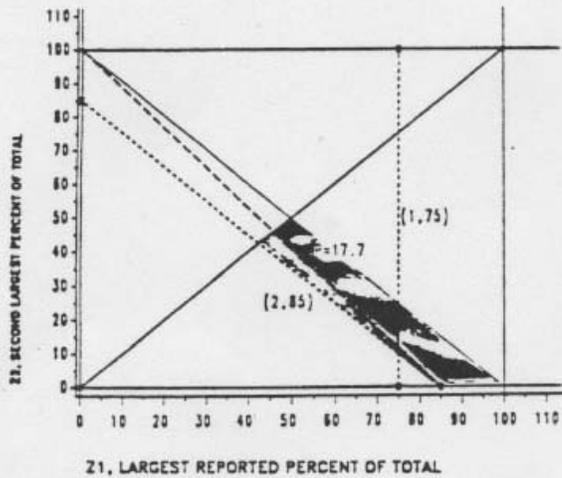
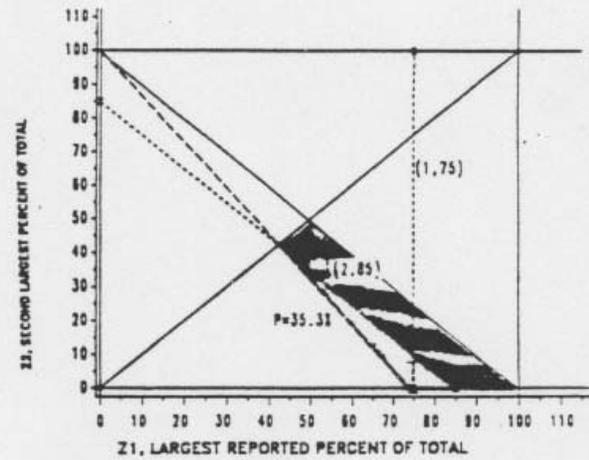


FIGURE 1D
 P-PERCENT MORE CONSERVATIVE THAN (2,85) AND (1,75)
 P = 35.3 PERCENT



Note: Values corresponding to cells in a table are in the triangle bounded by the lines $z_1 = z_2$, $z_2 = 0$, and $z_1 + z_2 = 1$. Values which correspond to sensitive cells are shaded.

Figure 1B shows the sensitivity regions for the p-percent rule with $p = 17.65$ and $p = 35.29$. The selection of these values of p will be discussed below. Note that if $p = 0$, the sensitivity line falls on top of the line $z_1 + z_2 = 1$. At that point there are no sensitive cells. Similarly if p is negative, there are no sensitive cells.

To Find p so that $S^{p\%}(\mathbf{X}) \leq S^{(n,k)}(\mathbf{X})$ for all cells, \mathbf{X} .

Consider the case where the (n, k) rule is being used and there is also a requirement that no respondent's contribution be estimable to within p-percent of its value. We would like to find the value of p so that the p-percent rule is closest to the (n, k) rule with $S^{(n,k)}(X) \geq S^{p\%}(X)$. Thus, there may be cells classified as sensitive by the (n, k) rule which would not be sensitive by the p-percent rule, but all cells classified as sensitive by the p-percent rule would be classified as sensitive by the (n, k) rule. Consider the $(2, 85)$ rule illustrated in Figure 1A. The p-percent rule, closest to the $(2, 85)$ rule, which would satisfy this requirement would be the one which intersects the line $z_2 = 0$ at the same point as the $(2, 85)$ rule. Thus, for a given value of k_2 we must have

$$S^{(n,k)}(X) \geq S^{p\%}(X)$$

Similarly, if we were first given the value of p for the p-percent rule, we must have

$$S^{(n,k)}(X) \geq S^{p\%}(X)$$

For the $(2, 85)$ rule, $p/100 = 15/85 = .1765$, so that $p = 17.65$ percent. Figure 1C shows the $(2,85)$ sensitivity region along with the less conservative $p = 17.65$ percent region.

For the $(1, k_1)$ rule, the p-percent rule closest to the $(1, 75)$ rule satisfying this requirement would be the one intersecting the line $z_1 = z_2$ at the point $(75, 75)$. For a given value of k_1 we must have

$$\frac{p}{100} = \frac{100}{k_1} - 2.$$

Similarly, if we were first given the value of p ,

$$k_1 = \frac{100}{\frac{p}{100} + 2}.$$

With $k_1 = 75$, the less conservative p-percent rule would have $p = -66.7$, which would result in no cell suppression. For $p = 17.65\%$, we would need $k_1 = 45.94$, a very restrictive rule.

To find parameter p so that $S^{p\%}(\mathbf{X}) \geq S^{(n,k)}(\mathbf{X})$ for all \mathbf{X} .

We would like to find the value of p so that the p-percent rule is closest to the (n, k) rule with $S^{k)(n)}(X) \leq S^{p\%}(X)$. Thus, there may be cells classified as sensitive by the p-percent rule which would not be sensitive by the (n, k) rule, but all cells classified as sensitive by the (n, k) rule would be classified as sensitive by the p-percent rule. Again, we consider the (2, 85) rule as illustrated in Figure 1A. In this case the most conservative p-percent rule needed would be the one that intersects the line $z_1 = z_2$ at the same point as the (2, 85) rule. Given the value of k_2 this leads to

$$\frac{p}{100} = \frac{200}{k_2} - 2.$$

If we were first given the value of p, we would need

$$k_2 = \frac{200}{\frac{p}{100} + 2}.$$

For $k_2 = 85$, this gives $p/100 = 200/85 - 2 = .3529$. Figure 1D shows the (2,85) sensitivity region along with the $p = 35.29$ percent region.

To find the most conservative p% rule needed to include the sensitivity region of the $(1, k_1)$ rule, we need the p-percent rule which intersects the line $z_2 = 0$ at the same point as the $(1, k_1)$ rule. Given the value of k_1 , this leads to

$$\frac{p}{100} = \frac{100}{k_1} - 1.$$

If we were first given the value of p, we would need

$$k_1 = \frac{100}{\frac{p}{100} + 1}.$$

For the (1,75) rule, this leads to $p/100 = 25/75 = .3333$.

To find the $(1, k_1)$ rule going through the same point as the $(2, 85)$ rule and the p -percent rule with $p = 35.29\%$, substitute the desired value of p into the above equation and find $k_1 = 73.91$.

In this case since we started with the $(2, 85)$ rule, which lead to $p = 35.29$, a consistently less conservative $(1, k_1)$ rule is the one that has $k_1 = 73.91$. Thus the p -percent rule with $p = 35.29$ provides slightly more protection than either the $(2, 85)$ rule or the $(1, 73.91)$ rule. Table 1 in the text summarizes these results for selected values of p , or equivalently for selected values of q/p .

Example

Consider the three cells below. Let x_1^k represent the largest value reported by a respondent in cell k ; x_2^k the second largest value reported by a respondent in cell k ; and so on. Here we assume that respondents report in only one of the cells 1, 2 or 3. Cell membership is denoted by the superscript k . Superscript T represents the total.

	Cell 1	Cell 2	Cell 3	Total
	$x_1^1 = 100$	$x_1^2 = 1$	$x_1^3 = 100$	$x_1^T = 100$
		$x_2^2 = 1$		$x_2^T = 100$
		$x_3^2 = 1$		$x_3^T = 100$
		.		.
		.		.
		.		.
		$x_{20}^2 = 1$		$x_{20}^T = 100$
SUM	100	20	100	220

Assume that we are using the (n, k) rule with $n = 2$ and $k = 85$ percent. As described above, the related rules are the p -percent rule with $p = 17.65$ (more conservative), the p -percent rule with $p = 35.29$ (less conservative) and the $(1, 73.91)$ rule.

Using any of these rules, Cell 1 and Cell 3 are clearly sensitive ($N = 1$, so $S(X) > 0$). It is also easy to verify that using any sensible rule Cell 2 is not sensitive. We consider two cells, the union of Cell 1 and Cell 2 and the Total.

The cell sensitivities for these rules are

$$\begin{aligned}
 S^{(2,85)}(\text{Cell 1} \cup \text{Cell 2}) &= 100 + 1 - 5.667 \cdot 19 = -6.67 \\
 S^{(17.6\%)}(\text{Cell 1} \cup \text{Cell 2}) &= 100 - 5.667 \cdot 19 = -7.67 \\
 S^{(1,73.91)}(\text{Cell 1} \cup \text{Cell 2}) &= 100 - 2.833 \cdot 20 = -43.34 \\
 S^{(35.29\%)}(\text{Cell 1} \cup \text{Cell 2}) &= 100 - 2.834 \cdot 19 = 46.16
 \end{aligned}$$

$$\begin{aligned}
S^{(2,85)}(\text{Total}) &= 100 + 100 - 5.667 \cdot 20 = 86.66 \\
S^{(17.6\%)}(\text{Total}) &= 100 - 5.667 \cdot 20 = -13.34 \\
S^{(1,73.91)}(\text{Total}) &= 100 - 2.833 \cdot 120 = -239.96 \\
S^{(35.29\%)}(\text{Total}) &= 100 - 2.834 \cdot 20 = 43.32
\end{aligned}$$

The union of Cell 1 and Cell 2 is not sensitive according to the (2, 85) rule and the 17.65% rule. However, both the (1, 75) and the 33.3% rule classify the cell as sensitive. Looking at the respondent level data, it is intuitively reasonable that the union of Cell 1 and Cell 2 is sensitive, even though the rule of choice for this example was to protect only against dominance by the 2 largest respondents. This cell corresponds to the point (83.3, .008) on Figure 1.

The Total is sensitive for the (2, 85) rule and the p-percent rule with p=35.3%. It is not sensitive for the (1, 73.9) rule or the p-percent rule with p=17.6%. This point corresponds with the point (45.5, 45.5) on Figure 1.

Consider the inconsistency in using the (2, 85) rule alone. In the above example, if the union of cell 1 and cell 2 (not sensitive by the (2, 85) rule,) is published, then the largest respondent knows that the other respondents' values sum to 20, and each of other respondents knows that the other respondents' values sum to 119. If the total (sensitive by the (2, 85) rule) is published then the largest two respondents each knows that the sum of the remaining respondents' values is 120, and each of the small respondents knows that the sum of the others' values is 219.

Intuitively, it would seem that more information about respondent's data is released by publishing the nonsensitive union of cell 1 and cell 2 than by publishing the sensitive total. The inconsistency can be resolved by using a combination of (n, k) rules, such as the (1, 73.91) and (2, 85), or by using a single p-percent rule with p = 35.29 or a pq-rule with q/p = 2.83. These changes result in additional, but more consistent suppressions.

Proponents of the simple (2, 85) rule claim that more protection is needed when respondents have competitors with values close to their own. Proponents of the simple (1, 75) rule claim that more protection is needed if the cell is dominated by a single respondent. These people argue that the use of a simple (n, k) rule allows them to determine which rules are needed for their special situations without the additional suppressions which would result from a more consistent approach.