AUTOMATED STATISTICAL DISCLOSURE CONTROL

Lawrence H. Cox, U. S. Bureau of the Census

THE SUPPRESSION PROBLEM

Statistical disclosure is said to occur when published data values may be combined to produce an inference about sensitive datum, or the identity or response of a subset of the respondents, which reveals "too much." In most applications, it is the confidentiality of the identity and responses of single respondents which must be maintained and the notion of revealing "too much" about a respondent is captured in the notion of sensitive cell. Typically, cell sensitivity is defined in terms of linear sensitivity measures, as discussed in [SA] and [CX3]. The n-respondent, k% dominance rule employed by the U.S. Bureau of the Census to define and identify potential breaches of respondent confidentiality in Economic Censuses is an example of an upper linear sensitivity measure. The (n,k) rule defines a cell (or cell union) to be sensitive if the aggregate of the response of n or fewer respondents in the cell exceeds k% of the total cell value. In general, $0 < n \leq 5$ and $50 \leq k \leq 100$.

Having established the definition of sensitive cell in a particular application, the disclosure practitioner next must quantify the notion of revealing "too much" about respondents in the cell. This amounts to establishing a definition of what constitutes an acceptable estimate of the value of a sensitive cell or cell union (i.e., an estimate of the cell value which can be tolerated) and what does not. This concept is generally defined in terms of an open interval (L(X),U(X)) containing the value V(X) of sensitive cell X. Estimates of V(X) which lie within or overlap this interval are defined as unacceptable while estimates of V(X) which strictly contain or do not meet this interval are by definition acceptable estimates of V(X). Ideally, to maintain the concept of statistical disclosure in a uniform manner between respondents and cells, the real numbers L(X) and U(X) should be defined for each X by formulae which are consistent with the operant sensitivity criterion.

Under a complete cell suppression methodology, all sensitive cells X must be suppressed, together with sufficiently many additional cells (called the complementary suppressions) to ensure that only acceptable estimates of the values of sensitive cells (or cell unions)[1] may be obtained by combining linear relationships between the suppressed cells. Ideally, complementary suppressions should be chosen so that, according to some defined measure of information, the information lost by suppressing the sensitive cells and the chosen set of complementary suppressions is no greater than that which would be lost if any other complete set of complementary suppressions had been chosen instead. As will be discussed in a later section of this paper, this ideal, though rarely attained, is crucial to controlling the suppression process theoretically and operationally.

The suppression problem has been described thus far as a problem of linear estimation of the values of sensitive cells. Complementary suppressions are made to prevent unacceptably narrow estimates of sensitive cells from being made; and a suppression pattern is considered to be complete when only acceptable estimates of the values of sensitive cells may be made through analysis of the pattern. These two principles come together when techniques of linear estimation are employed to derive inferrable ("actual") estimates of sensitive cells, so that these inferrable estimates may be compared with predetermined acceptable estimates of these cells. If all actual estimates are acceptable, then the suppression pattern is by definition complete. Otherwise, if at least one actual estimate violates the defined acceptable limits of protection, then additional complementary suppression must be performed.

Two principles emerge from the above characterization. The first and more obvious of these is that more accurate techniques of linear estimation may uncover more subtle cases of disclosure and should therefore result in increased protection of respondent confidentiality. The second principle is that, as the central problem in suppression methodology is that of establishing a balance between maintaining respondent confidentiality and the responsibility to publish as much meaningful data as possible to meet the legitimate information needs of the user community, then oversuppression of data can best be minimized by employing cell suppression techniques which are driven by and can operate in tandem with the techniques of linear estimation being employed. These dual problems of linear estimation and cell suppression are the main topics of this paper.

TECHNIQUES OF LINEAR ESTIMATION

The problem of determining optimal actual upper and lower estimates of the values of suppressed cells or cell unions in a publication system such as a census or survey may be structured mathematically as a linear optimization problem over a system of linear constraints, as follows. Each cell to be published for a particular statistic is assigned its value and each cell whose value is to be suppressed for this statistic is replaced by a unique variable in all linear relationships between the cell values. The linear estimation problem for each suppressed cell is then that of maximizing and minimizing the value of its corresponding variable to determine the smallest interval in which the value of this variable may vary. The linear estimation problem may be divided into two parts. First, we must construct a spanning set of linear equations for the linear space of all linear equations between the cell values derivable from the disaggregation relationships between the cells (i.e., a spanning set for the linear space of the values of all unions and differences of the suppressed cells which are effectively published). Second, optimal or adequate techniques of linear estimation must be applied to the variables corresponding to the sensitive cells in these equations to determine intervals in which their values may vary. Moreover, if, for the purpose of analysis and suppression strategy, this linear space is to be partitioned into subsets, a proper order in which to analyze these subsets must be determined so that a well-defined procession of linear estimation-suppression-linear estimation may be invoked.

The problem of determining a spanning set for this linear space is solved by embedding the set of disaggregation relationships between the cells and their corresponding linear equations within the structure of a Boolean lattice. This is accomplished by analyzing the definitional units which delimit the cells (such as various geographic parameters, age, income, or, for a business establishment, type of business operation) to identify generic classes of disaggregation relationships which exist between the cells. For example, if, for a particular statistical item, statistical values at the state level are disaggregated by the values corresponding to the constituent counties of each state, then all such disaggregation relationships and their corresponding linear equations form one class of disaggregation relationships (equations). Each such equation is represented by a node on a lattice and two nodes are defined to be hierarchically related in the lattice if and only if the equation corresponding to the second represents a disaggregation of a variable which appears as a constituent variable in the disaggregation represented by the first equation. For example, an equation which represents the disaggregation of a statistical value for a particular state by its constituent counties is hierarchically related to (above) an equation which represents the disaggregation of one of these constituent counties by its county-parts. All other relationships between the lattice nodes are defined by transitivity, so that the ambient lattice is Boolean. The interested reader is referred to [CX3] for a more complete description of the lattice procedure.

The lattice represents the most detailed aggregation-disaggregation schema for the system of cells and the linear relationships between their values derivable from the cell structure. Working from the atomic nodes (the "bottom") of the lattice (respectively, the maximal nodes of the lattice) and proceeding upwards (respectively, downwards), one may iteratively construct each aggregation (respectively, disaggregation) relationship between the cells and its corresponding linear equation. In particular, the set of all linear equations defined at the lattice nodes generates the linear space of all derivable linear equations between the cell values, thereby solving the first problem previously stated. Moreover, the lattice induces a natural ordering on the nodes, so that the lattice at the same time provides a means of partitioning the entire system of linear equations into a collection of subsystems (namely, into the sets of equations corresponding to the disaggregation of a single statistical item at each node) and a natural ordering of these subsystems for linear and suppression analysis.

A global solution to the problem of linear estimation of the values of suppressed cells may be obtained through application of techniques of linear programming. The linear equations (identified above) which span the linear space of all linear relationships between the cell variables form a constraint system on these variables. To determine optimal actual upper and lower estimates of any variable, say $\underline{x}$, it suffices to maximize (respectively, minimize) the objective function OBJFUN = x subject to these constraints. In most large applications, the number of variables and equations involved renders the global linear programming approach impractical or impossible to implement. However, if in a particular application it is computationally feasible to employ linear programming globally and if all actual estimates thus obtained for the values of the sensitive cells are acceptable, then the entire system is globally disclosure-free. However, if any actual estimate of the value of some sensitive cell is unacceptable, then further suppression is required (i.e., more variables must be introduced into the system). At this point, the strategy of applying a general linear programming package to the disclosure problem exhibits serious drawbacks, as it does not address the following important questions. Which additional cell(s) should be suppressed to protect this sensitive cell? Do the new suppressions adequately protect the sensitive cells, or do they also require protection? Does there exist a smaller set of covering suppressions which would have also protected the sensitive cells, so that fewer cells may be suppressed in the system, thereby minimizing oversuppression of data? A general purpose linear programming package, even one with the requisite speed and sophistication to completely analyze a large disclosure problem globally, provides only best linear estimates of prespecified objective functions and no indication of the structure of the disclosure problem from which sound complementary suppression strategies could be developed. Therefore, mathematical and mathematical programming techniques must be developed and applied selectively to a particular disclosure problem so that as much relevant complementary suppression information as possible is developed during linear estimation.

In large disclosure applications, such as a census or major survey, it is computationally infeasible to attempt a global solution to the estimation-suppression problem. In such applications, techniques such as the lattice procedure must be employed to partition the global problem into a sequence of well-defined local problems. Optimal or adequate techniques of estimation-suppression may then be employed locally to each subproblem in turn. If global optimality is sought, separation or decomposition results must be developed to ensure that global optima are also achieved locally.

The lattice procedure partitions the global problem into a sequence of local problems which may be represented in tabular form. Each tabular array represents the disaggregation of a particular statistical cell by as many classes of subcells as there are independent definitional units into which the cell may be disaggregated. For example, if the "parent" cell contains all retail sales establishments in a particular state, and if data are to be published for this state by type of business enterprise and by county, then these relationships may be represented by a two-dimensional tabular array, each internal cell of which represents the total of the values of the particular statistical item for all business establishments of a given type within a particular county. If the tabulated data were further broken down by type of sales activity, these tabular arrays would be three-dimensional.

In the discussion to follow on techniques of linear estimation, we shall restrict our attention to the problem of linear estimation in simple tables, i.e., in two-dimensional tabular arrays in which the grand total entry is unsuppressed. Although no one of the techniques we shall describe is totally limited in its applicability to simple tables, as simple tables possess a great deal of mathematical structure, they provide a natural context in which to analyze and compare different techniques of linear estimation.

The problem of determining optimal estimates of suppressed entries in simple tables is a transshipment problem as studied in the field of operations research. We refer the reader to [DN] for a discussion of transportation and transshipment problems, and limit this discussion to a characterization of their salient properties.

As an optimization problem, the transshipment problem seeks to minimize a given linear objective function $\sum_{i=1}^{r} c_i x_i$ of the variables x_i, subject to linear constraints among the x_i described by the matrix equation $AX = B$, where $X = (x_1 \; x_2 \; \ldots\ldots x_r)^t$. The mathematical property which characterizes transshipment problems is that each column in the constraint matrix A contains precisely two non-zero entries. One may readily observe that the linear estimation problem for simple tables (the problem of maximizing and minimizing the x_i's subject to the row and column equations in the table) is also of this form by writing down its constraint matrix A which consists of one column for each suppressed entry and one row for each of the row and column equations between the suppressed entries in the simple table. Each row of A is a vector of 0's, $+1$'s, and -1's (the entry is 0 if the corresponding variable does not appear in row or column equation in the simple table corresponding to this row of A, $+1$ if the variable corresponds to an internal entry in the simple table and does appear in the corresponding equation, and -1 if the variable corresponds to a marginal entry in the simple table and appears in the corresponding equation). As each entry appears in precisely two of these equations (i.e., one in its row and one in its column in the table), then each column of A contains precisely two non-zero entries. In general, all entries in the simple table are not suppressed, so that there are fewer variables than entries. As the values corresponding to non-variable entries in the table may be subtracted from their corresponding row and column marginal total entries, they may be considered to be "effectively zero." In general, therefore, the linear estimation problem for simple tables is a transshipment (or transportation) problem with zero-restricted variables (see [DN]).

This characterization of the linear estimation problem for simple tables is of considerable importance theoretically, as it brings with it proven techniques for computing these optima against which new techniques may be evaluated. If the location pattern of the zero-restricted variables is favorable, then the linear estimation problem can be solved computationally by transshipment theory techniques (such as the least-cost rule) in a generally efficient manner. However, there is no guarantee that this situation will be maintained from one table to the next. Moreover, straightforward application of transshipment theory software to the problem has the same drawbacks as previously described for the general linear programming approach if additional suppression is necessary. Techniques of linear estimation, no matter how computationally quick, cannot be truly efficient in the disclosure setting if they do not provide informative feedback to the complementary suppression process.

In [SA], a method is described for obtaining all effectively published aggregations with non-negative coefficients of the suppressed entries, and the observation that this set of aggregations forms a convex cone in the first orthant is made. From the theory of convex cones, it is well-known that this cone may be described as the set of all non-negative linear combinations of its extremal rays. The extremal rays are finite in number and correspond to the set of all elementary aggregations of the system, i.e., the minimal generating set over the non-negative coefficients of the set of all effectively published non-negative aggregations of the suppressed entries. Therefore, techniques such as the Chernikova algorithm may then be employed to compute the extremal rays of this cone (see [CH]). Upper estimates of non-negative aggregations of the suppressed entries may then be obtained by examining non-negative linear combinations of minimal sets of elementary aggregations which cover the given aggregation. Cox [CX4] showed that this procedure is exact, i.e., it does in fact produce optimal upper estimates of the values of such aggregations, as an effectively published minimal covering aggregation which achieves the optimum can always be constructed by adding only non-basic variables to the given aggregation.

Also in [CX4], a procedure for applying the convex cone procedure for achieving maxima of non-negative aggregations in conjunction with techniques of interval arithmetic to generate minima of such aggregations is developed. The procedure is as follows. Given a non-negative aggregation Z of the suppressed entries, one may generate all effectively published minimal covering aggregations of Z of the form Z+Y=a by the convex cone technique described above. For each such Y, the same technique may be employed to determine b=Max(Y), the maximum value Y may attain. The lower estimate Z $\geq$ a-b then follows. By exploiting properties of non-basic variables, Cox [CX4] shows that Min(Z), the minimum value of Z, must be attained by this procedure.

Since the cone-interval arithmetic technique of linear estimation explicitly constructs effectively published aggregations, it is capable of providing relevant information directly to the suppression process. This is a tremendous advantage. However, when the number of elementary aggregations is large, the technique of employing the convex cone-interval arithmetic approach to generate minima of only even single suppressed entries will exhibit serious computational drawbacks due to the complexity and number of the covering aggregations to be examined.

Previously in this paper, we discouraged the application of a general purpose linear programming package to the linear estimation problem for two principal reasons. In terms of computational feasibility and efficiency, even if the global problem is relatively small or if it is partitioned into subproblems, such as single tabular arrays, to which a linear programming package may be applied individually, this approach may not be computationally efficient or cost-effective as many pivots and optima must be generated. Also, these techniques of linear estimation provide the suppression process with minimal information. However, techniques of linear programming may be applied in a selective manner locally to eliminate these shortcomings. We briefly sketch this procedure.

The simplex algorithm employs a well-defined sequence of Gaussian eliminations ("pivots") to generate the optimum of a linear objective function subject to linear constraints by successively examining certain vertices of the feasible region. These vertices correspond to effectively published aggregations of the variables in which the basic variables achieve their optimum values, (see [DN] for a detailed discussion of the linear programming problem and the simplex algorithm). When a "black-box" linear programming package is employed, the user simply provides a single objective function with the constraint equations and receives the desired optimum as output, so that the iterations of the simplex algorithm are transparent to the user. In the disclosure setting, however, it is precisely the set of these intermediate aggregations that are of interest as they may yield optima of other aggregations of interest. Moreover, whenever such an aggregation violates the acceptable estimates of a particular cell, it must be made available to the suppression process for analysis and augmentation. Therefore, it is desirable to devise a computationally efficient procedure which will capture these intermediate aggregations. What follows is a description of a rudimentary version of such a procedure.

The simplex algorithm minimizes the value of a given objective function (aggregation) as follows. Beginning with a representation of the linear constraints in feasible canonical form and a linear aggregation (objective function) Z in the non-negative variables x, effectively published aggregations (linear combinations of the constraint equations) are successively added to the relation OBJFUN=Z until a linear equation of the form, B = OBJFUN-c, is obtained, where B is a non-negative linear combination of non-basic variables (and hence achieves zero as its minimum value) and c is constant. The conclusion c = Min(Z) is therefore reached. The reader will note that this is equivalent to deriving the equation Z-B=c with B and c as above. Analogously, the simplex algorithm may be employed to maximize the aggregation Z by applying the above techniques to minimize OBJFUN = (-Z), so that a relationship of the form Z+B'=a may be derived, with B' a non-negative linear combination of non-basic variables and a constant, so that a = Max(Z).

Each change of basis (pivot) operation in the simplex algorithm generates a new basic feasible solution of the linear program. If the objective function is kept fixed and degeneracy of the linear program is avoided, then each pivot further reduces the value of the objective function until optimality is reached. However, the observations stated in the preceding paragraph can be applied after each pivot operation to identify maxima and minima of aggregations other than the given objective function. We illustrate this procedure for the case of obtaining optima of single variable objective functions (i.e., single suppressed entries) in algorithmic form.

Simplex Technique for Linear Estimation of Suppressed Entries

0. Place the system of constraint equations in feasible canonical form.

1. For the NEXT variable x both of whose optima have not been found, set the objective function OBJFUN = x (if the minimum value of x is sought) or OBJFUN = -x (if the maximum value of x is sought). If all x have been optimized, END.

2. Examine this basis row by row to see if it optimizes any basic variable y both of whose optima have not already been determined, as follows. If a row has not changed since the last pivot, it is not examined.

   a. Let the NEXT row of the tableau to be examined contain the basic variable y (with coefficient +1).

If this row contains only non-negative entries, then the constant value a of the effectively published aggregation corresponding to this row is equal to the maximum value y may attain. If there are no other non-zero entries on this row, then the value of y is effectively published, so that Min(y) = a also and we may GOTO the NEXT row to be examined.

Moreover, if all entries on this row are non-negative and if the maximum value b of the corresponding linear combination of the non-basic variables is known, then a-b is equal to the minimum value of y. This technique is particularly useful when this linear combination consists of one non-basic variable whose maximum has already been determined.

If Max(y) and Min(y) were attained in this step, GOTO c.

b. If the only non-negative entry on this row is the +1 coefficient of y, then the constant value of c corresponding to this row is equal to the minimum value of y.

Moreover, if all entries on this row except that corresponding to y have non-positive coefficients, and if the maximum value of d of the negative of the corresponding linear combination of the non-basic variables is known, then the maximum value of y equals c+d. Again, this technique is most useful when this linear combination consists of the negative of a single non-basic variable whose maximum has already been determined.

c. The minimum value of any non-basic variable z must equal zero.

3. If this basis optimizes OBJFUN, GOTO 1.

4. Pivot according to the simplex algorithm to improve this non-optimal basic feasible solution for the objective function OBJFUN.

5. GOTO 2.

This algorithm is quite rudimentary and could certainly benefit from improvements indicated by theoretical insight or practical experience. The choice of the method employed to achieve feasible canonical form will certainly affect the performance; and the efficiency of the algorithm should vary considerably with the choice of the NEXT variable to be analyzed. No optimizing principle currently exists for these choices. Nor is it clear that an objective function must remain fixed until optimality is reached - perhaps a change of objective function or a brief digression in the analysis to optimize a new objective function before returning to the original tableau would be worthwhile. Also, if estimates are to be made of virtually arbitrary sets of aggregations, then a sharpening of this technique is necessary. These topics are under investigation. Preliminary tests of this algorithm on simple tables are most encouraging. Optima of all suppressed entries in many test tables were obtained with less than one-third as many pivots as variables once feasible canonical form had been achieved.

COMPLEMENTARY CELL SUPPRESSION

In the preceding section, a great deal of emphasis was placed upon capturing those effectively published aggregations which give rise to unacceptable estimates of suppressed cells, in order to make these aggregations readily available to the complementary suppression process. Assuming that this has been accomplished, the methodological problem which is immediately posed is that of how to best utilize this information. Specifically, can a complementary suppression methodology be developed which takes full advantage of the information provided by these effectively published aggregations?

At the crux of the suppression problem is the problem of adequately protecting the sensitive cells while minimizing the adverse impact of the suppression process on the quality and information content of the published data. The information content of a set of publication cells may be measured in numerous ways. Many of these measures are in some way dependent upon the number of cells suppressed. The principle of "not suppressing more cells than is necessary" seems to be widely accepted. In general, therefore, the count of the number of suppressed cells is a relevant measure of the adverse impact of the disclosure process on the information content of the published data. Accordingly, we henceforth assume that this principle obtains, and say that one suppression pattern is inferior to a second if both patterns adequately protect the sensitive cells and if the second pattern involves fewer complementary suppressions than the first. For purpose of this discussion, we define oversuppression as the suppression of more than the minimum number of cells necessary to adequately protect the sensitive cells, and thereby implicitly assume that the cells are of equal weight in terms of their importance for publication. We further limit the discussion to the determination of optimal or adequate patterns of complementary suppression in single tabular arrays.

Even granting the above restrictions, and assuming that the operant sensitivity criterion admits the existence of corresponding subadditive and superadditive upper and lower sensitivity measures, few theoretical results exist for the problem of minimizing the number of complementary suppressions necessary to render an arbitrary tabular array disclosure-free. This is an area which requires greater research activity, particularly in three and higher dimensions.

In one dimension, the problem is trivial. As there is only one equation corresponding to the one-dimensional tabular array, then the subadditivity and superadditivity of the sensitivity measures imply that the optimal upper estimate of a suppressed internal entry is equal to the difference between the marginal entry and the total of the unsuppressed internal entries in the array if the marginal entry is unsuppressed; and is equal to an optimal upper estimate of the marginal entry if the marginal entry is suppressed. If the marginal entry is suppressed, its optimal lower estimate is equal to the sum of the unsuppressed internal entries. Optimal lower estimates of suppressed internal entries are equal to zero; and optimal upper estimates of a suppressed marginal entry must be determined from other tables. Therefore, to protect a suppressed internal entry from above (respectively, a suppressed marginal entry from below), in a one-dimensional tabular array it suffices to suppress sufficiently many internal cells so that the sum of their values equals or exceeds the amount of additional suppression required to adequately protect the original internal (marginal) suppression from above (below).

Extending this technique directly to two-dimensional tabular arrays (tables) results in tremendous shortcomings. Effectively, only line estimates are being considered, and no account is taken of the effect of a complementary suppression made in, say, a particular row has upon its column. A suppression strategy which takes advantage of the transshipment problem attributes of the disclosure problem in two-dimensional tabular arrays is necessary to produce minimal suppression patterns. Unfortunately, no such strategy exists.

To bring the salient mathematical properties of the problem to light, we make certain uniformizing assumptions. Specifically, we assume that any disclosed row or column in the table can be rendered disclosure-free once a total of two suppressions (counting the sensitive cells) have been made on the row or column. Under this assumption, the following theorem of Cox from [CX1] completely solves the two-dimensional suppression problem for line estimates of suppressed entries, and is hence a partial solution to the two-dimensional suppression problem.

THEOREM. Let R denote the number of disclosed rows (i.e., rows containing one suppression which must be protected) and C denote the number of disclosed columns in a two-dimensional tabular array. Assume that any one additional suppression in each disclosed row or column will suffice to render this row or column disclosure-free. If R = C = 1 then at most three (3) additional suppressions are necessary to render all rows and columns in the table disclosure-free. Otherwise, Max (R,C) additional suppressions suffice.

The proof of this theorem in the general case is constructive and proceeds as follows. For definiteness, assume R = Max (R,C). Then the first C complementary suppressions are to be chosen in each of the C disclosed columns, provided that each is chosen in a different disclosed row. The remaining (R-C) disclosed rows may be protected in any manner whatever (i.e., by any (R-C) complementary suppressions each in one of these disclosed rows), provided that if one such suppression is chosen in a column containing no suppressions, then at least one other such suppression is chosen in this column as well. It results that the number of such minimal suppression patterns is a combinatorial function of Max (R,C) which grows like the factorial, so that many such patterns may be identified in general. Extensive testing experience with live data by the U.S. Bureau of the Census of an automated complementary disclosure system for two-dimensional tabular arrays based upon this theorem has borne out its practical value. The system applies the theorem to attempt to identify an optimal pattern involving only one additional suppression in each disclosed row and column. If the resulting pattern is not complete, the theorem is applied repeatedly until a complete pattern is achieved.

As previously stated, this theorem is only a partial solution to the two-dimensional suppression problem as, within the confines of its hypotheses, it only protects sensitive cells in their rows and columns by insuring that each row and column aggregation of suppressed entries contains at least two variables, and does so in the minimum amount of complementary suppression, but does not insure that all effectively published aggregations of suppressed entries contain at least two variables. A corresponding theorem which for a given suppression pattern would identify all minimal covering suppression patterns in which each effectively published aggregation contains at least two variables would, within our limiting assumptions, completely solve the two-dimensional suppression problem. If such a theorem does not exist (in the sense that, in general, the minimum number of complementary suppressions necessary to complete a given suppression pattern cannot be expressed in close form), then a computationally efficient algorithm for constructing such minimal completing patterns (perhaps one based upon branch and bound techniques) must be discovered.

Virtually nothing is known theoretically about the suppression problem in three and higher dimensions. The two-dimensional theorem stated above may be expressed in terms of the König-Egerváry Theorem[2] which is a statement about bipartite graphs. To indicate the difficulty of generalizing the two-dimensional suppression theorem to three-dimensions, we remark that no generalization of the König-Egerváry Theorem to tripartite graphs exists.

References and footnotes, omitted due to space considerations, are available from the author.