Journal of Official Statistics, Vol. 16, No. 3, 2000, pp. 243–253

The General Application of Significance Editing

David Lawrence[1] *and Richard McKenzie*[2]

The significance editing technique reduces and prioritises the input editing workload in a sample survey. The method is based on estimating the impact of resolving individual edit queries on the survey estimates. Using these impacts it is possible to ignore queries of trivial value while making sure that the most serious queries are given the most attention. The result is efficient allocation of resources to editing without compromising data quality. Lawrence and McDavitt (1994) described an implementation of significance editing in the Australian Survey of Average Weekly Earnings. This article extends the previous work by describing the general principles underlying the significance editing technique, and offering suggestions for applying the technique to other surveys.

Key words: Data quality; editing; respondent burden; respondent recontact; survey costs.

1. Introduction

A commonly used output editing tool in sample surveys is to break down an estimate into contributions from individual survey respondents and then examine the data from the largest contributors. In most surveys it is necessary to have the full complement of data available to calculate these individual contributions, so output editing tools like this are usually run near the end of the processing cycle. Significance editing is based on a similar concept, but applied to the input editing phase. It extends the usual edit rules defined for the input editing phase by calculating the impact on the final estimates of amending the queried data to bring it into line with our expectations based on the edit rules. In most cases it is possible to estimate these impacts without referring to any other reported data in the current processing cycle other than that from the respondent currently being processed.

Lawrence and McDavitt (1994) described the implementation of significance editing in the Australian Survey of Average Weekly Earnings (ASAWE). The principles that underpin the significance editing application in the ASAWE can be applied to a wide range of surveys, including surveys that are much more complicated than the ASAWE. This article describes the general principles underlying the significance editing technique, and gives suggestions for their application in other surveys.

In recent years there has been a marked shift in the attitude towards editing. The desirability of correcting as many errors in the survey returns as possible has been rethought in the light of resource constraints and to encompass quality improvement in the total

[1] Centre for Health Services Research, Department of Public Health, University of Western Australia, Nedlands, WA 6907, Australia. E-mail: davidl@dph.uwa.edu.au
[2] Statistical Support Unit, Australian Bureau of Statistics, GPO Box K881, Perth, WA 6001, Australia.

survey process (Linacre and Trewin (1989), Lepp and Linacre (1993)). Granquist (1995) describes the evolution in attitudes towards data editing. Granquist (1997) proposes a new view on editing with an emphasis on minimising the resources expended on editing, targetting editing effort where gains in quality can be made, and that the primary focus of editing should be on identifying error sources to prevent errors from arising (particularly in ongoing surveys). This concept of selective editing, where not all survey returns that are flagged for attention during the editing phase are actively pursued, involves focussing the editing effort on errors of greatest influence on the final data. Several methods have been suggested for targetting the editing, including the ''score variable'' method of Hidiroglou and Berthelot (1986), the weighted deviance from group medians of Van de Pol (1997), and the methods of scoring edit failures of Latouche and Berthelot (1992) and Lawrence and McDavitt (1994).

Significance editing allows the resources allocated to input editing to be reduced and prioritised. It can also reduce respondent burden. By estimating the effect on the survey estimates of resolving the edit queries for each unit, it is possible to prioritise the edit queries for follow-up action. It also allows an assessment of the impact of editing, and gives useful information for determining the amount of editing to be performed. The most significant queries can be given the most attention, and those of trivial importance can be ignored altogether. This can allow an improvement in overall quality of the final estimates if the resources saved are used for a more extensive follow-up of the most important queries by highly skilled and trained editors and other quality improvement activities.

In this article, the term *editing* refers to the process by which survey returns are checked for possible errors. This process includes recontacting respondents where suspect data is identified and the resulting amendment of data where errors are found. Editing can be regarded as being split into two phases – *input editing* and *output editing* (Hughes et al. 1990). The input editing phase involves checking individual survey returns in isolation from other responses in the current processing cycle. Typically the input editing procedure involves applying a set of edit rules, or edits, which flag questionable responses for follow-up attention. The respondents can then be recontacted to either verify or amend the reported data. Output editing involves identifying possible errors by examining all of the survey returns for the current cycle as a whole.

2. Controlling the Size of Editing Errors

For many years the goal of editing has generally been seen as identifying all, or as many as possible, of the errors that can be found in the survey returns. A key thrust of significance editing is that this is an unnecessary and unrealistic goal. Results from sample surveys are subject to sampling error and a range of nonsampling errors. Nonsampling errors are difficult to measure, and most survey results are published with indications of the sampling errors only. For this reason, it is important to keep the extent of the nonsampling errors small, relative to the sampling error. However, it is not efficient to expend a lot of resources on attempting to reduce respondent reporting errors and data entry errors to zero. The size of the ''editing bias'' due to not checking all returns should definitely by controlled. However, once it is reduced to less than say 10–20 per cent of the relevant sampling error (and other sources of error), further reductions have no real impact on the inferences that are made from the survey results.

There is ample evidence to show that in many cases over-editing of survey returns occurs. Granquist and Kovar (1997) provide a detailed review of research in this area. For instance, Pullum, Harpham and Ozsever (1986) reported that in the World Fertility Survey the editing had little noticeable effect on the estimates or the analysis of the data other than delaying the release of the results by about a year. Latouche and Berthelot (1992) reported on a trial application of a process similar to significance editing in the Canadian Annual Retail Trade Survey. They showed that respondent recontacts could be restricted to units with higher scores with only minor effects on the survey aggregates. Lawrence and McDavitt (1994) cut editing in the ASAWE by about half without any noticeable effect on the survey results. Linacre and Trewin (1989) found that most changes made in the Australian manufacturing survey had little effect and that the majority of the total change due to editing was due to a few changes to large units. Boucher (1991) found that the number of edit rules applied in the Canadian annual survey of manufactures could be reduced to just a few rudimentary edits with little difference in the final data.

The idea of significance editing is to only edit survey returns where there is some benefit in doing so. The amount of editing to be performed in a collection can be set based on reasonable accuracy requirements for the final survey outputs.

3. Elements of the Significance Editing Strategy

The fundamental principle of significance editing is to estimate the impact of editing a survey return, and use that measure of impact to prioritise any edit query action. Note that significance editing extends the edit rules that are defined for the survey, and is dependent on well-designed edit rules to be successful. Significance editing will not improve on a poor set of edit rules or enable the identification of errors that are not addressed by any edit rules. It is still necessary to define logical and comprehensive edit rules according to reasonable expectations of the population being surveyed.

The key elements of the significance editing strategy are as follows:

- choosing the expected amended values;
- determining local scores;
- combining local scores to produce a global score;
- setting cut-off scores.

The process of editing a survey return is then to apply the edits to each return; score returns that fail any edits; determine the global score and compare the score to a pre-determined cut-off value. Follow-up action is undertaken if the global score exceeds the cut-off. However, some edits may remain outside of the significance editing process. For instance you may choose to exclude fatal edits (e.g., ensuring additivity of data items that sum to a total) and always follow up failures of these edits. It seems sensible that if a unit is judged significant for follow-up action, all edit failures would be examined.

Each of the elements of the significance editing strategy are described in more detail below.

4. Choosing the Expected Amended Values

Most types of edits can be thought of in terms of a predetermined view of how the population should behave. We will refer to this as an editing model. In some cases this model might be quite simple. For instance, an edit rule may require that a set of components add up to a total, or that complementary variables correspond (e.g., there must have been at least one employee working overtime if overtime allowances have been paid). More complex edits may be based on historical comparisons, or expectations based on the behaviour of other, similar units. Whatever the type of edit rule, the editing model implicit in the edit rule can be used to generate an *expected amended value*. This is a value that is more likely, according to the editing model, than the actual reported value.

The general principle is to choose an expected amended value that is in line with the editing model being used. Determining the expected amended values relies on subject matter knowledge, in the same way that developing the edit rules does. Determining a full set of expected amended values can be a time-consuming task. Indeed it may involve weeks of effort for a complex survey with many data items, although for repeated surveys this is a one-off cost, and the cost is certainly small compared to the resources expended on editing.

The following procedures could be used for handling some different types of edit rules:

– In a range edit that tests if a reported value lies inside a range (a, b), the mid-point of the range could be chosen as the expected amended value. This type of edit is commonly used for comparing data with responses in a previous cycle of an ongoing collection.
– For a logical edit which checks that two data items together have sensible responses (e.g., total wheat harvest compared to area sown to wheat) the expected amended value would be a value that makes the reported data most logical (e.g., if a wheat harvest was reported without any wheat crop sown, the expected amended value for wheat crop sown could be the area you would expect to sow to produce a harvest of the reported size).
– For a ''cross-survey'' edit that compares similar (or identical) data items reported by the same respondent on different surveys, choose the value corresponding to the other survey as the expected amended value.
– Suppose an edit is designed to detect if respondents are ''cheating'' in their returns in an ongoing survey by always supplying the same figures when such consistency is unreasonable (the so-called ''inlier'' problem, see Granquist and Kovar (1997)). For such an edit, the expected amended value would be chosen based on the average or expected size of typical fluctuations in similar units.
– For data that may not always be reported, zero may be a sensible expected amended value. For example, in an edit testing if termination payments or directors' fees are more than a certain percentage of the gross wages and salaries bill, an expected amended value of 0% would be a conservative choice.

Where two or more variables are being compared by an edit you can generate expected amended values for each variable, and use each to determine the significance of an amendment. For example, in an edit examining the average of two data items an expected amended value may be generated for either the numerator or the denominator. In many cases the results will be very similar either way, and it is a matter of judgement to decide which of the variables should be given an expected amended value for the calculations.

For the purpose of significance editing, it is not necessary that the expected amended value be estimated sufficiently accurately to be usable as an imputed value. It is merely used to gain a comparison, across units, of the relative importance of resolving each edit query. Indeed it is quite possible for a unit to fail an edit because its data is atypical according to the editing model, even though the data may be correct. It is valid to edit such a unit, as you would want to confirm the correctness of such atypical data, but it would be inappropriate to use the expected amended value as an imputed value in estimation.

5. Calculating the Local Score

Once an expected amended value has been generated, it can be used to estimate the likely impact on the overall estimate of amending the data for the queried unit. The principle here is to estimate the change in the estimate that would occur if the unit under consideration had its data changed to the expected amended value(s), leaving all other units intact. The expected change in the estimate for a particular variable that would come about from amending one particular unit's responses is called the *local score* for that variable. Suppose you have an estimate of the total of the form $\hat{X} = \sum w_i x_i$, where w_i is the weight for unit i, and x_i is the reported value for unit i. For data item X, the local score would then be given by

$$S_i = w_i \, \delta x_i$$

where δx_i is the difference between the reported value and the expected amended value for the ith unit.

The score needs to be evaluated as the editing is being done, so it should be independent of other units in that processing cycle. This may require approximation in some cases. For instance survey weights are often adjusted for nonresponse, but the final nonresponse rate is not known until all the data has been processed. The design weight, based on an assumption of 100% response, can be used to calculate the local scores.

For the purposes of significance editing, the score only needs to be a relative measure. The absolute value is not important. So using the design weight causes no problems provided the nonresponse occurs at random.

More complex estimators would result in different score functions. Suppose you are estimating a more complex statistic that is a function of several variables. Let $\hat{X} = f\{\bar{x}_{rh}\}$, where $r = 1, \ldots, R$ represents R separate variables, and $h = 1, \ldots, L$ are strata. Then by expansion in Taylor series to first order, the approximate score is given by

$$S_i = \sum_{r=1}^{R} \frac{\partial f}{\partial \bar{x}_{rh}} \delta \bar{x}_{rh} \tag{1}$$

where unit i is in stratum h, the partial derivative is evaluated at the point of the expected value of $\bar{x}_{rh}$, and $\delta \bar{x}_{rh}$ represents the change in the mean of variable x_r in stratum h as a result of replacing the reported value for unit i with the expected amended value.

For example, in the ASAWE the estimates are ratios of the form $\hat{Z}/\hat{X}$ where $\hat{Z} = \sum w_i z_i$ is an estimate of total earnings and $\hat{X} = \sum w_i x_i$ is an estimate of total employment. Using (1), the score for a unit i would be given by

$$S_i = \frac{1}{\hat{X}} \left(w_i \, \delta z_i - \frac{\hat{Z}}{\hat{X}} w_i \, \delta x_i \right)$$

This first order approximation compares very closely to the actual score function for this example given by Lawrence and McDavitt (1994). This methodology can be used to generate a score function for many common estimators. Note that it makes sense to specify an expected amended value for the numerator or the denominator, but not necessarily both, so one of δz_i or δx_i could be zero.

As can be seen in this example, evaluation of the partial derivatives may require knowledge of the final survey estimates. Again, the score is used as a relative measure to assign editing priorities. The absolute value is not critical. The ASAWE simulation study found that the score function is not particularly sensitive to the specification of these values (as they are fixed for all units and the key importance of the score is the relative values) and thus estimates from a previous cycle can be used for an ongoing survey, or reasonable expectations for a one-off survey.

The expected amended value does not need to be estimated with high precision. It is not part of significance editing to use the expected amended value as an imputed value. The principle is to choose a value that is in line with the editing rule being tested, which may not be close to the true value for every unit. The ASAWE simulation study found that the process was not very sensitive to the specification of the expected amended values, with a $\pm 20\%$ change in expected amended values having negligible effect on the outcome.

6. Combining Local Scores to Determine a Global Score

There can be many edit rules applied to each survey return, and it is possible that one return will be flagged for follow-up action by more than one rule. However, respondent recontact needs to be performed on a unit by unit basis, not separately for each data item. This can be achieved by having a procedure to combine the local scores into one global score for the unit. Lawrence and McDavitt (1994) chose to take the maximum absolute value of the local scores as the global score. This principle ensures that a significant error for any variable will receive attention. In other surveys, an issue in determining the global score will be scaling variables that use different units of measurement or naturally have a different order of magnitude. The suggestion is that one of the following be used:

– The absolute pseudo-bias of Latouche and Berthelot (1992)

$$\text{Scaled score} = \frac{S_i}{|\hat{X}|}$$

– The relative pseudo-bias of Lawrence and McDavitt (1994)

$$\text{Scaled score} = \frac{S_i}{SE(\hat{X})}$$

The relative pseudo-bias has the advantage of linking the size of the editing errors to another measure of errors – the sampling error. If large nonsampling errors are known to be associated with particular variables, then this information can also be incorporated into determining the scaled local score. For instance, if respondent record-keeping practices are such that only an approximate figure is ever available for a particular data item, there may be little point editing it highly, even if the sampling error was very small. In this situation the scaled score can be based on the mean squared error:

$$\text{Scaled score} = \frac{S_i}{MSE(\hat{X})}$$

Once the local scores have been appropriately scaled, the global score can be taken as the maximum of the absolute values of the scaled local scores. This global score can then be used to prioritise the unit for follow-up action.

7. Setting Cut-off Scores Using a Simulation Study

Once global scores are determined, a cut-off score should be set so that survey returns with a score above the cut-off are queried. Priority for edit follow-up action is then determined by the relative size of the global score, with the greatest attention being focussed on survey returns with the highest scores. The cut-off scores need to be set to control the impact of editing errors on the survey estimates. This is most easily done in an ongoing collection. The procedure involves capturing the raw data (as provided on the original survey return before any query action takes place) and the final survey data after the usual input editing occurs.

The simulation study then involves retrospectively scoring the raw data (using the expected amended values generated from the editing model, not the final survey data after editing). The survey returns are then ranked by the global score. Choose an arbitrary percentage p. Then create a composite survey returns file taking the final survey data for the top $p\%$ of survey returns, and the raw survey data for the remainder. Estimates are recomputed from this composite file. This procedure is repeated for different values of p, building up a relationship between the amount of editing performed and the effect on the estimates. Cut-off scores can then be chosen so that the bias due to not editing some of the survey returns is low compared to other sources of error in the survey such as the sampling error.

The simulation study approach can also be used to evaluate the effectiveness of the editing model, by calculating the percentage of edit queries that generate amendments in the data. This may indicate a problem with particular edits, for instance tolerances being set too tight. It is also possible to identify if units that receive high scores in the significance editing process are the units most likely to receive amendments at the recontact stage. This is one way of measuring the effectiveness of the procedures used to assign the expected amended values.

Setting cut-off scores by means of a simulation study is relatively simple. It is also quite powerful in terms of providing objective evidence to users of the survey data that the significance editing strategy will be efficient without compromising the quality of the final data.

8. Setting Cut-off Scores Without a Simulation Study

The methodology of the simulation study is simple and effective and should be repeatable in most periodic surveys. It may not always be practical to conduct a simulation study, however. Obviously it is not possible for one-off surveys, and if the survey is run infrequently, the time taken to collect sufficient study data may be too long. Also there may be some cases where limits in the survey processing system will preclude the collection of the original unedited survey responses.

In these cases it may still be possible to determine cut-off scores using a type of response model. The following is an example of a simple model that may be used. More complex models could also be derived for specific situations.

In order to avoid making assumptions about the distribution of survey data items and reporting errors, we will use the significance editing score S_i as an estimate of the impact on the survey estimate of resolving the edit queries for the ith unit. If a cut-off a is set for determining when follow-ups will occur, then for those units not edited the S_i will lie in the range $(-a, a)$. Suppose that $E(S_i) = 0$. Assuming that the editing errors are independent between units that triggered edits, the bias in the survey estimate of not editing a set of n_a survey returns can be approximated by the sum of the S_i values over those n_a unedited survey returns. It is simple to show that the squared bias due to not editing a set of n_a survey returns is given by:

$$E(bias^2(\hat{X})) = n_a E(S_i^2)$$

To ensure that inferences from the survey results remain unaffected by not editing some of the survey returns, it is important to keep this bias low compared to other sources of error in the survey, such as the sampling error. Assuming that the S_i are distributed uniformly on $(-a, a)$ (which should be quite a conservative assumption), we can say $E(S_i^2) = a^2/3$. If we want to set the cut-off a to achieve $E(bias^2(\hat{X})) \leq k\, Var(\hat{X})$ for some $k < 1$, it can be seen that a should be set so that

$$a \leq \sqrt{\frac{3k}{n_a}} SE(\hat{X})$$

If we set the cut-off as

$$a = \sqrt{\frac{3k}{n}} SE(\hat{X})$$

then, this will give

$$E(bias^2(\hat{X})) \leq k\, Var(\hat{X})$$

With the number of conservative assumptions involved in the above model, this may give quite low cut-off scores. To evaluate this model we revisited the data of Lawrence and McDavitt (1994). The predicted cut-offs under this approach were calculated for a value of $k = 0.1$, and these were compared with the actual cut-off values in use in the survey in Table 1. As there are several variables in the ASAWE we chose to evaluate the model using the variable with the highest standard error.

The model suggests that higher cut-offs could be set in the ASAWE than those that are being used if we were aiming to restrict the bias to this level, which would result in a further reduction in editing. At the time of original implementation of significance editing in the ASAWE we knew the cut-offs were being set conservatively in terms of their effects on the survey estimates, but the cut-offs were quite radical in terms of the reduction in editing effort involved. The amount of editing performed was reduced by more than

Table 1. Significance editing cut-off scores used in ASAWE, compared with cut-off scores predicted by model, by state

State	Model cut-off	Cut-off used
New South Wales	0.11	0.04
Victoria	0.18	0.06
Queensland	0.20	0.11
South Australia	0.28	0.15
Western Australia	0.26	0.11
Tasmania	0.40	0.31
Northern Territory	0.41	0.33
Australian Capital Territory	0.44	0.33

half; however, the results of the original simulation study, and the results of our simple model, suggest further cuts could be made.

There are drawbacks to using a model-based approach rather than conducting a simulation study. For instance, there is the question of the assumption made in the model. If errors are systematically biased in one direction the effect of significance editing on the estimates could be much more marked. In such cases, the bias in the errors would need to be built into the model or, preferably, the simulation study approach be employed. Our recommended approach is to use quality control procedures to establish significance editing in a collection where prior information is not available for performing a full simulation study. In this approach at the start of processing the original cut-off scores can be set by making conservative assumptions as to the distribution of the S_i values. Then as processing gets underway records can be kept of the original survey data as reported by respondents, and the amended values that result from performing the edits. These data can then be used to conduct a simulation study to refine the cut-off scores.

9. Staff Acceptance of Significance Editing

One of the major obstacles to implementing significance editing in an ongoing collection is the cultural shift associated with the underlying attitude to editing. In many collections a large proportion of the available resources is expended on editing. There has been a long tradition of regarding the role of editing in assuring data quality as being vitally important. The notion that some errors can be left uncorrected is counter-intuitive to many people. It can raise the question of why respondents are selected in the first place if their data are not considered significant. Of course, it is not the data itself that is insignificant, but the effect of pursuing the edit queries when balanced against other uses of the available processing resources.

To this end we found the conduct of the simulation study to be very useful. While cut-off scores can be set using a response model, justifying the choice of the model is best achieved through something like our simulation study. The results of the simulation study for the ASAWE were stark and dramatic. Once the results had been presented many people who had initially been sceptical as to whether significance editing could work took a much more favourable attitude to the technique.

It is also valuable for processing staff to understand why it is important to collect data from a respondent when errors in that data can be accepted. Particularly when sampling from skewed populations it is unavoidable that some respondents will have a higher influence on the overall estimates than others, and of course some respondent errors will be larger than others.

One issue relating to the implementation of significance editing is the important issue of the loss of jobs in editing if editing is cut. Significance editing is not, per se, a tool for reducing the amount of editing performed in a survey. We suggest that the amount of editing in a collection should be set based on some reasonably objective criteria, to achieve control over the size of expected bias due to respondent reporting errors, and to ensure that this component of error is kept small in comparison to other sources of error in the survey. In the ASAWE editing resources were cut because the previous editing strategy was inefficient and targeted far too many units for follow-up action. This need not have been the case, although we suspect that many surveys will have taken a cautious approach to their editing strategies and will be over-editing.

Another important aspect of significance editing is the prioritisation of edit queries that it provides. While units whose global scores fall below the cut-off can be ignored for editing, those with higher scores can be given relative priority in editing according to the size of their global score. Many editing strategies in current use do not provide such a method of prioritisation. It is a naive view to assume that recontacting the respondent elicits the correct survey response. We know from practical experience that the amount of probing and the quality of probing during recontact can influence the outcome dramatically. Some respondents will confirm their reported data with plausible justifications at the first query; but with more extensive probing from highly trained and experienced editors, problems can be detected and resolved. As noted by Granquist (1997), the primary focus of editing should be on identifying error sources to prevent errors arising, particularly in ongoing surveys, and the resources saved from not recontacting some units can be redirected towards addressing the most serious edit queries in greater depth. The outcome of implementing significance editing may not be cost cutting but rather an improvement in overall data quality.

Another good opportunity arising from significance editing is addressing other sources of nonsampling error. Nonsampling errors are usually difficult to measure, and in many surveys editing of survey returns is the only regular activity addressing nonsampling errors. Where resources can be saved from editing as a result of a significance editing process, these can be redirected towards other quality initiatives. These could include addressing coverage problems on sampling frames, examining questionnaire design, running post-enumeration studies, and so forth. The overall result can be an improvement in the quality of the survey data for the same amount of resources expended.

10. Conclusion

Significance editing is a method for establishing a more logical and rigorous footing for designing an editing strategy. It addresses the questions of how much editing should be done and what tolerances should be set for the edits. The procedure is based on controlling the size of residual bias due to not resolving all edit queries to meet predetermined quality targets. The procedure also gives a relative priority to each unit targeted for follow-up action, so that the most important queries can be pursued more rigorously. The result is a more rigorous justification of the resources spent on editing, and the respondent burden imposed by the edit recontacts.

Significance editing brings some traditional output editing concepts into the framework of input editing. By focusing on the impact of editing on the final analysis, the editing strategy can be more closely linked to the overall accuracy of the final outputs from the survey.

11. References

Boucher, L. (1991). Micro-editing for the Annual Survey of Manufactures: What is Value Added? Proceedings of the 1991 Annual Research Conference, U.S. Bureau of the Census, 765–781.

Granquist, L. (1995). Improving the Traditional Editing Process. In Business Survey Methods, eds. B.G. Cox, D.A. Binder, B.N. Chinnappa, A. Christianson, M.J. Colledge, and P.S. Kott, New York: Wiley, 385–401.

Granquist, L. (1997). The New View on Editing. International Statistical Review, 65, 381–387.

Granquist, L. and Kovar, J.G. (1997). Editing of Survey Data: How Much Is Enough? In Survey Measurement and Process Quality, eds. L. Lyberg, P. Biemer, M. Collins, E. de Leeuw, C. Dippo, N. Schwarz, and D. Trewin, New York: Wiley, 415–435.

Hidiroglou, M.A. and Berthelot, J.M. (1986). Statistical Editing and Imputation for Periodic Business Surveys. Survey Methodology, 12, 73–84.

Hughes, P.J., McDermid, I., and Linacre, S.J. (1990). The Use of Graphical Methods in Editing. Proceedings of the 1990 Annual Research Conference, U.S. Bureau of the Census, 538–550.

Latouche, M. and Berthelot, J.M. (1992). Use of a Score Function to Prioritize and Limit Recontacts in Editing Business Surveys. Journal of Official Statistics, 8, 389–400.

Lawrence, D. and McDavitt, C. (1994). Significance Editing in the Australian Survey of Average Weekly Earnings. Journal of Official Statistics, 10, 437–447.

Lepp, H. and Linacre, S.J. (1993). Improving the Efficiency and Effectiveness of Editing in a Statistical Agency. Bulletin of the International Statistical Institute. Proceedings of the 48th Session, Florence, 111–112.

Linacre, S.J. and Trewin, D.J. (1989). Evaluation of Errors and Appropriate Resource Allocation in Economic Collections. Proceedings of the 1989 Annual Research Conference, U.S. Bureau of the Census, 197–210.

Pullum, T.W., Harpham, T., and Ozsever, N. (1986). The Machine Editing of Large Sample Surveys: The Experience of the World Fertility Survey. International Statistical Review, 54, 311–326.

Van de Pol, F. (1997). Selective Editing in the Netherlands Annual Construction Survey. In Statistical Data Editing (Volume 2); Methods and Techniques, United Nations, New York.

Received January 1999
Revised April 2000