

SERIES

STUDIES AND PERSPECTIVES

96

ECLAC SUBREGIONAL
HEADQUARTERS
FOR THE CARIBBEAN

Statistical disclosure control for Caribbean census tables

**A proposal to expand
the availability
of disaggregated
census data**

Francis Jones



UNITED NATIONS

ECLAC

Thank you for your interest in this ECLAC publication



Please register if you would like to receive information on our editorial products and activities. When you register, you may specify your particular areas of interest and you will gain access to our products in other formats.

 www.cepal.org/en/publications



www.cepal.org/apps

SERIES

STUDIES AND PERSPECTIVES

96

ECLAC SUBREGIONAL
HEADQUARTERS
FOR THE CARIBBEAN

Statistical disclosure control for Caribbean census tables

A proposal to expand the availability
of disaggregated census data

Francis Jones



E C L A C

This document has been prepared by Francis Jones, Population Affairs Officer of the Statistics and Social Development Unit of the Economic Commission for Latin America and the Caribbean (ECLAC) subregional headquarters for the Caribbean. The study benefited from the comments received from the following reviewers: Jean-Louis Tambay, Researcher of Statistics Canada; Ari Silva, Consultant of the REDATAM Development Centre, Latin American and Caribbean Demographic Centre (CELADE) - Population Division of ECLAC; Nyasha Skerrette and Machel Pantin, both Economic Research Assistant of the ECLAC subregional headquarters for the Caribbean.

The views expressed in this document, which has been reproduced without formal editing, are those of the author and do not necessarily reflect the views of the Organization.

United Nations publication
ISSN: 1728-5445 (electronic version)
ISSN: 1727-9917 (print version)
LC/TS.2020/159
LC/CAR/TS.2020/6
Distribution: L
Copyright © United Nations, 2020
All rights reserved
Printed at United Nations, Santiago
S.20-00895

This publication should be cited as: F. Jones, "Statistical disclosure control for Caribbean census tables: a proposal to expand the availability of disaggregated census data", *Studies and Perspectives series-ECLAC Subregional Headquarters for the Caribbean*, No. 96 (LC/TS.2020/159-LC/CAR/TS.2020/6), Santiago, Economic Commission for Latin America and the Caribbean (ECLAC), 2020.

Applications for authorization to reproduce this work in whole or in part should be sent to the Economic Commission for Latin America and the Caribbean (ECLAC), Publications and Web Services Division, publicaciones.cepal@un.org. Member States and their governmental institutions may reproduce this work without prior authorization, but are requested to mention the source and to inform ECLAC of such reproduction.

Contents

Abstract.....	5
Introduction	7
I. Statistical disclosure control for census tables	9
A. The problem of statistical disclosure control for tabular census data.....	9
B. Current practice in the Caribbean.....	12
II. A review of SDC methods for census tables and international best practice.....	15
A. Pre-tabular methods	15
B. Post-tabular methods.....	16
1. Small cell adjustments.....	17
2. Random rounding.....	17
3. ABS cell perturbation (or the cell key method).....	19
C. Differential privacy	23
D. Applicability of SDC methods to Caribbean censuses	24
III. Comparative analysis and testing of methods.....	27
A. Comparative analysis of methods.....	27
B. Testing of selected post-tabular methods	31
1. Measurement of disclosure risk and data utility	31
2. A description of the test data.....	32
3. Results of comparative testing	34
4. An additional feature of the cell key method	36
IV. Implementation of disclosure control functionality in redatam.....	39
A. Post-tabular perturbation (with cell keys).....	39
B. Complementary methods.....	40
C. Disclosure control protection for continuous variables	40
D. REDATAM's programming window.....	41

V. Conclusions	43
Bibliography.....	45
Annexes	49
Annex 1	50
Annex 2	52
Studies and Perspectives-The Caribbean Series: issues published.	55

Tables

Table 1	A census table from which identification (or self-identification) could occur	10
Table 2	An example of individual attribute disclosure	10
Table 3	An example of group attribute disclosure	11
Table 4	An example of disclosure by differencing: activity status by area for all persons	12
Table 5	An example of disclosure by differencing: activity status by area for all persons without a disability	12
Table 6	Part of a Statistics Canada census table: population by age and sex, Census 2016	18
Table 7	Random rounding to base five: rounding probabilities, Statistics Canada.....	18
Table 8	ABS cell perturbation applied to a cross-tabulation of religion by district (Belize census, 2010)	21
Table 9	Cross-tabulations of census data and disclosure risk.....	33
Table 10	Measures of disclosure risk and information loss for post-tabular disclosure control methods applied to two-, three-, four- and five-way cross-tabulations of data from the Belize census, 2010	35
Table 11	Lookup table (or 'ptable') for the perturbation scheme CP1.....	37
Table 12	Perturbation of totals and marginal totals using the cell key: impact on disclosure risk for cross-tabulations of data from the Saint Kitts and Nevis census, 2011.....	37
Table A1	Measures of disclosure risk and information loss for post-tabular disclosure control methods applied to two-, three-, four- and five-way cross-tabulations of data from the Jamaica census, 2011	50
Table A2	Measures of disclosure risk and information loss for post-tabular disclosure control methods applied to two-, three-, four- and five-way cross-tabulations of data from the Saint Kitts and Nevis census, 2011.....	51
Table A3	Lookup table (or 'ptable') equivalent to transition matrix B in diagram 1 (first 13 rows only)	52

Figures

Figure 1	Post-tabular perturbation schemes for census data (the symmetric case)	30
----------	--	----

Diagrams

Diagram 1	Two alternative transition matrices for cell perturbation	28
-----------	---	----

Abstract

Among Caribbean statistical offices, there is widespread awareness of the issue of statistical confidentiality which has often led to quite strong restrictions on the availability of detailed and disaggregated census tables. In North America, Europe, Australia and New Zealand, statistical offices use disclosure control methods, particularly methods involving some form of data perturbation, to safely publish detailed and disaggregated census tables. These methods are not used in the Caribbean (or Latin America) but could safely facilitate the release of much more detailed census tables than has hitherto been possible in many Caribbean countries. This study reviews the problem of statistical disclosure control for census tables and international best practice, focusing particularly on the use of perturbative methods; it carries out comparative analysis and testing of the cell perturbation method and random rounding methods; and recommends that these methods should be made available to statistical offices through the REDATAM software.

Introduction

Statistical disclosure control has been an area of increasing concern for official statisticians over recent decades. The increased use and availability of official statistics, the richness of the data environment into which official statistics are now published, and the development of modern computing and data science have all increased the threats to statistical confidentiality and therefore the importance of statistical disclosure control. This concern and attention to the issue of disclosure control has been most evident among the statistical offices and statistical systems of North America, Europe, Australia and New Zealand. Among Caribbean statistical offices¹ there is both an awareness and concern about the issue which influences the way in which statistics are published and disseminated, although relatively little use has been made of specific methods for disclosure control, especially methods involving data perturbation.

The problem of statistical disclosure control for census data has received particular attention because complete enumeration has important implications for disclosure risk. It is relatively easier to identify an individual in a census dataset and to learn some information about them, compared to a survey dataset. This is because in a census dataset it is known, at least with high probability, that all members of the population are present in the data. A household survey, by contrast, only contains a small fraction of the population and this uncertainty about who is, or is not, in the sample provides some protection against the risk of disclosure. At the same time, the main purpose of a census is to be able to provide information about small geographic areas and rare population characteristics which would not be captured accurately in a survey. It is precisely these data for small geographic areas and rare population characteristics which are most likely to lead to disclosure of information about individuals. In small Caribbean countries, census datasets do not have to be disaggregated very far, by geography and other variables, before disclosure risks become very apparent.

¹ For the purposes of this analysis, Caribbean refers to the following countries and territories: Anguilla, Antigua and Barbuda, Aruba, Bahamas, Barbados, Belize, Bermuda, British Virgin Islands, Cayman Islands, Curaçao, Dominica, French Guyana, Grenada, Guadeloupe, Guyana, Jamaica, Martinique, Montserrat, Saint Kitts and Nevis, Saint Lucia, Saint Vincent and the Grenadines, Sint Maarten, Suriname, Trinidad and Tobago, Turks and Caicos Islands and United States Virgin Islands.

Census frequency tables for small areas and sub-populations will commonly contain small cell counts. This leads to the possibility that individuals or households could be identified and information about them disclosed. Such disclosure of confidential information would constitute not just a breach of one of the United Nation's fundamental principles of official statistics but would also, in many countries, contravene national statistical laws. At present, statistical offices seek to prevent this from occurring by avoiding publication of tables which contain small cell counts. This is commonly achieved by, for example, not releasing data for small geographic areas and redesigning tables to avoid presenting information on rare population characteristics. This will clearly reduce the usefulness of census data outputs for some users at least.

There is extensive literature addressing the problem of statistical disclosure control (SDC) for population and housing censuses. Methods, policies and practices have been developed to enable statistical offices to maximise user access to census data while reducing to an acceptable minimum the risk that there is any breach of statistical confidentiality. This is referred to as the risk-utility trade-off. There are two broad categories of disclosure control methods. Non-perturbative methods generally involve data suppression, either through restricting the publication of certain tables or table redesign. Perturbative methods, on the other hand, change rather than suppress data, introducing some form of random variation or noise (see Antal, Enderle and Giessing, 2017). The statistical disclosure control literature subdivides into two main strands: disclosure control for tabular data and disclosure control for microdata. Disclosure control for census microdata is not considered here but was addressed in a relatively recent ECLAC study on the dissemination of Caribbean census microdata (Jones and Fox, 2016).

The motivation behind this study arises primarily from ECLAC's work with Caribbean statistical offices, supporting them in the online dissemination of census data through REDATAM. There are currently ten Caribbean countries offering access to one or more census datasets through an online query system, using REDATAM. In the design of these applications, disclosure control was a major factor which shaped the design of the application and determined what data were released or not released.

This study makes the case that for Caribbean census data, a better trade-off between risk and utility can be achieved by applying statistical disclosure control methods which have not previously been applied in Latin American or Caribbean censuses, specifically perturbative methods. It also proposes that implementation of these methods in the REDATAM software would be the most effective way to make them available to Caribbean statistical offices and other REDATAM users.

I. Statistical disclosure control for census tables

A. The problem of statistical disclosure control for tabular census data

Tables of frequency counts are the primary form in which most census data are published. They are cross-tabulations of two or more variables, with each cell containing a count of the number of persons (or households, dwellings or buildings) falling within that cross-classification of the variable categories.² Disclosure risks arise primarily when there are table cells containing only a very small number of cases, because those combinations of variable categories are rare in the population. This is most especially true for cells containing only one case, that is, a unique combination of variable categories. Such cases are at risk of being identified, meaning that someone recognises that the 1 in the cell belongs to a known individual or household. Cells containing counts of 2 are also regarded as risky since if one person can be matched to the cell containing two cases, then the other individual is effectively unique by subtraction within the cell.

The literature on statistical disclosure control formalises this problem of disclosure in a slightly more precise way (see Hundepool and others, 2012 or Hundepool and others, 2010), distinguishing between several different types of disclosure. The simplest is identification, simply that an intruder is able to link a data item to one known individual (or household). For example in table 1, identification would be said to have occurred if somebody knowing either a Hispanic Methodist or a White Seventh Day Adventist was able to recognise that one of the counts of 1 in either of those two cells must belong to one of those individuals. In order to establish a link with a known person, the intruder must already know the individual's ethnicity and religion, and therefore the table would not reveal any new information about those individuals to the intruder, except the fact that they are unique in the population (which could be significant). For this reason, identification is regarded as a somewhat less serious form of disclosure, but it is still a concern, both because of the perception that individuals are visible in the data and because identification can be the first step towards actually disclosing information about an individual. There may be other tables where this uniqueness can be exploited to reveal information.

² The REDATAM software (like SPSS) distinguishes between frequency tables, which show frequency counts for one variable, and cross-tabulations which show frequency counts for cross-classifications of two or more variables. In common with most of the statistical disclosure control literature, this study uses the term frequency table to include both one dimensional and multi-dimensional tables of frequency counts.

Table 1
A census table from which identification (or self-identification) could occur
(Numbers of people)

	Anglican	Baptist	Methodist	Pentecostal	Catholic	Seventh Day Adventist	Total
Black	7 284	2 347	7 126	4 572	1 836	2 417	25 582
East Indian	20	58	9	123	40	10	260
Hispanic	7	5	1	120	171	53	357
Mixed	138	39	93	157	309	65	801
White	153	45	51	13	351	1	614
Other	11	15	6	9	80	10	131
Total	7 613	2 509	7 286	4 994	2 787	2 556	27 745

Note: Table is based on artificial data for illustration purposes only.

A more serious form of disclosure is individual attribute disclosure. This occurs when not only is an individual identified, but information about them is revealed. Table 2 shows a three-dimensional cross-tabulation: ethnicity by religion by some health condition. If in a cross-tabulation of three variables, there are individuals who can be identified uniquely by just two variables, then for those individuals, their values on the third variable will be revealed. In this case, the Hispanic Methodist does suffer from the health condition, but the White Seventh Day Adventist does not.

Table 2
An example of individual attribute disclosure
(Numbers of people)

Ethnicity	Health condition	Anglican	Baptist	Methodist	Pentecostal	Catholic	Seventh Day Adventist	Total
Black	No	6 795	2 268	6 616	4 408	1 746	2 342	24 175
	Yes	489	79	510	164	90	75	1 407
	Total	7 284	2 347	7 126	4 572	1 836	2 417	25 582
East Indian	No	18	55	7	115	40	10	245
	Yes	2	3	2	8	0	0	15
	Total	20	58	9	123	40	10	260
Hispanic	No	7	5	0	116	166	50	344
	Yes	0	0	1	4	5	3	13
	Total	7	5	1	120	171	53	357
Mixed	No	131	37	89	157	289	64	767
	Yes	7	2	4	0	20	1	34
	Total	138	39	93	157	309	65	801
White	No	148	45	49	12	345	1	600
	Yes	5	0	2	1	6	0	14
	Total	153	45	51	13	351	1	614
Total	No	7 099	2 410	6 761	4 808	2 586	2 467	26 131
	Yes	503	84	519	177	121	79	1 483
	Total	7 602	2 494	7 280	4 985	2 707	2 546	27 614

Note: Table is based on artificial data for illustration purposes only.

In fact, it is not always necessary for identification to take place, for attribute disclosure to occur. Table 3 shows another three-dimensional cross-tabulation, this time occupation, by age, by health insurance status. The presence of zeros discloses the fact that all managers aged 15-24 have health insurance but among agricultural workers aged over 65, none of them have health insurance. This is referred to as group attribute disclosure.

Disclosure does not have to result from the appearance of small cell counts in tables. The difference between two or more larger numbers can imply the existence of a small cell count and so lead to disclosure. Tables 4 and 5 could very plausibly be produced using a flexible table generator like those built using REDATAM.

They show activity status by area for all persons (table 4) and activity status by area for persons without a disability (table 5). Although there are no small cell counts in either table 4 or 5, differencing the two tables reveals the existence of six persons with a disability that have a unique combination of activity status and area. This example of disclosure by differencing illustrates how disclosure control for census tables is not simply a matter of addressing the small cell counts, but also the implied small cell counts which can be derived by differencing.

Table 3
An example of group attribute disclosure
(Numbers of people)

Occupation	Health insurance	15-24	25-34	35-44	45-54	55-64	65+	Total
Managers	No	0	5	14	36	26	14	95
	Yes	4	10	58	129	123	872	1 196
	Total	4	15	72	165	149	886	1 291
Professional/ Technical	No	8	17	36	81	55	22	219
	Yes	16	39	140	290	148	112	745
	Total	24	56	176	371	203	134	964
Clerical	No	6	14	18	40	21	4	103
	Yes	19	43	55	117	53	7	294
	Total	25	57	73	157	74	11	397
Service/sales	No	29	40	49	56	31	9	300
	Yes	30	38	37	52	23	5	180
	Total	59	78	86	108	54	14	480
Agriculture	No	6	30	34	36	31	5	142
	Yes	2	16	22	15	4	0	59
	Total	8	46	56	51	35	5	201
Total	No	49	106	151	249	164	140	859
	Yes	71	146	312	603	351	991	2 474
	Total	120	252	463	852	515	1 131	3 333

Note: Table is based on artificial data for illustration purposes only.

It should also be remembered that although the focus here has been on individuals, tables of data about households, dwellings or buildings can also lead to disclosure of information, including about the individuals belonging to those statistical units. Indeed, there are particular disclosure risks associated with variables such as household type, size, income level or dwelling type and with variables that tend to be shared by household members, for example ethnicity and religion. Household variables, by definition, group individuals and if households can be identified, this also constitutes disclosure of information about the household members. Individual variables such as religion or ethnicity, particularly those which are uncommon in a given area, could lead to a cell count of, say, 5 corresponding to the members of one household. Attention needs to be paid to these hierarchical patterns in the way that variables are made available through an online query system, particularly the way that building, dwelling, household and individual variables can be used in combination.

These examples of potential disclosure might seem to suggest that it is just a small matter relating to a few unusual cases. If census tables are only distributed through a single national census report, this might be true. Online dissemination, however, particularly through flexible table generating applications, dramatically increases the ability of statistical offices to provide detailed access to census data. A REDATAM application can make thousands of times more tables available to users than a census report. It is this fuller and richer exploitation of census data that inevitably comes into conflict much more directly with the obligation to protect statistical confidentiality. The ability of users to define their own tables raises the possibility that if an individual is identified, it may be possible to recover their entire census record. Furthermore, recent work on database reconstruction attacks shows that the reconstruction of census micro-datasets from tabular outputs, at least in part, is much more feasible than was previously thought (see chapter II, section C).

Table 4
An example of disclosure by differencing: activity status by area for all persons
(Numbers of people)

	Area 1	Area 2	Area 3	Area 4	Area 5	Area 6	Total
Worked	4 923	658	887	1 367	1 520	1 834	11 189
Unemployed	538	58	104	247	175	185	1 307
Home duties	430	65	70	93	112	116	886
Attended school	1 125	257	91	173	216	534	2 396
Retired	854	97	106	148	181	250	1 636
Unable to work	139	8	31	65	37	45	325
Total	8 009	1 143	1 289	2 093	2 241	2 964	17 739

Note: Table is based on artificial data for illustration purposes only.

Table 5
An example of disclosure by differencing: activity status by area for all persons without a disability
(Numbers of people)

	Area 1	Area 2	Area 3	Area 4	Area 5	Area 6	Total
Worked	4 871	650	883	1 346	1 507	1 817	11 074
Unemployed	535	56	104	237	169	184	1 285
Home duties	419	64	63	87	107	112	852
Attended school	1 120	256	90	173	215	533	2 386
Retired	742	93	81	122	157	217	1 412
Unable to work	72	5	14	30	16	24	161
Total	7 759	1 124	1 235	1 994	2 171	2 887	17 170

Note: Table is based on artificial data for illustration purposes only.

The undertakings given to census respondents about protecting the confidentiality of the information they provide are unambiguous. Statistical offices take their responsibilities seriously in this regard and work hard to prevent worst-case disclosure scenarios and the damage they could do to institutional reputations and trust in official statistics. Yet statistical offices also have a responsibility to maximise the utilisation of census data and to provide the richest and most detailed data products that they can to census users, which leads back to the risk that information about individuals is disclosed. Statistical disclosure control is about finding the best compromise between these conflicting obligations.

B. Current practice in the Caribbean

In producing census tables, such as those appearing in national census reports or other similar publications, statistical offices in the Caribbean generally take care to limit or avoid the publication of tables with small cell counts that could lead to disclosure. They restrict the publication of tables for small geographic areas and small population subgroups, and tables may be redesigned to prevent disclosure, for example by recoding variables. In the absence of dedicated capacity in the area of statistical disclosure control, subjective decision making and the attitudes of individual statisticians towards statistical confidentiality and managing risk inevitably play a major role.

This kind of table suppression or redesign is considered a form of statistical disclosure control but, used alone, it does not strike an efficient balance between disclosure risk and data utility. It is certainly possible to reduce disclosure risk to a low level, but only by heavily restricting the publication of statistics about small geographic areas, small population subgroups and detailed breakdowns of variables like occupation for which there is a real user demand.

Where REDATAM is used to disseminate census data, disclosure risk is managed through the design of the online application. Since the volume of tables available to the user through an interactive tool is much greater, this becomes an even more complex problem. Relevant aspects of design include the variables which are available in the application, the way in which those variables are coded, the extent to which users can

cross-tabulate different variables against each other (or use variables to filter database queries), the level of geographic disaggregation which is available, and whether the user is provided with functionality to use REDATAM's programming language to specify database queries.

REDATAM is not a single online table generator, rather it is a tool that statistical offices can use to build flexible table generators. Many statistical offices have worked with ECLAC, or sometimes independently, to provide online access to their census data using REDATAM. Statistical disclosure control itself is as much an art as a science and this, combined with the different attitudes to statistical confidentiality in different countries, means that there are variations from country to country in the way that REDATAM applications have been designed and the implications of those designs for disclosure risk.

Whatever decisions statistical offices have taken about the publication of census tables, relying solely on non-perturbative methods is not an effective strategy. This is true both for the release of standard census tables and for the design of REDATAM applications. The result is either excessive disclosure risk or heavy restrictions on the release of census data and suppression of data. Perturbative methods introduce the possibility that for a very small loss of data accuracy, statistical offices could significantly expand the publication of census data, safely releasing data that would otherwise remain unpublished.³ A range of perturbative methods have been used by statistical offices around the world to achieve more effective disclosure control for census data and the following chapter of this report will review some of these methods.

³ The same logic also applies to statistical disclosure control for microdata where, although the methods are rather different, the most effective anonymization of microdata is normally achieved through a combination of non-perturbative and perturbative methods.

II. A review of SDC methods for census tables and international best practice

A range of statistical disclosure control methods have been proposed and applied to protect tabular data, but only a more limited selection of methods have been deemed appropriate for census use. Not only do censuses tend to generate large numbers of relatively complex tables but, as mentioned in the previous chapter, complete enumeration of the population makes the census something of a special case. Some methods which are quite widely applied to other statistical outputs are generally deemed inappropriate for use with census data. The cell suppression method involves the primary suppression of risky cells and the secondary suppression of non-risky cells (in order to prevent primarily suppressed values being re-calculated from the values published in the table). However, cell suppression is not normally applied in censuses because of the relative difficulty of applying primary and secondary suppressions across a large number of linked tables (Shlomo, 2018). The controlled rounding method protects tabular data through rounding while maintaining additivity between rounded interior cells and the rounded total and marginal totals, but it depends on computationally expensive linear programming which makes it unsuited to the size, scope and magnitude of census tables (Shlomo, 2018).

This chapter looks at three approaches to statistical disclosure control for census tables: pre-tabular methods, post-tabular methods and differential privacy, and considers their applicability to census data in the Caribbean. Each of these approaches involve introducing, in different ways, random variation into the data in order to reduce to an acceptable minimum the risk either that an individual could be identified in a table or that information could be disclosed about an individual. Clearly it is desirable that the random variation introduced to the data is the minimum that is necessary to achieve this objective.

A. Pre-tabular methods

The most widely used method of pre-tabular perturbation is record swapping which has been used in recent censuses in the United States (censuses 1990, 2000 and 2010), United Kingdom (censuses 2001 and 2011), Sweden (2011), Austria (2011) and Belgium (2011). The essence of the record swapping (or data swapping)

method is that a small proportion of households which are similar in terms of some basic social and demographic characteristics are swapped, that is, their geographic codes are swapped as if the households physically swapped places. The swapping routine can be refined to favour households with rare characteristics which are at greater risk of identification or attribute disclosure (this is referred to as targeted record swapping). The geographical distance between swapped households can also be varied according to disclosure risk, within a certain limit (a household would not be swapped with another in a completely different part of the county).

The rationale for swapping the geographic location of households is that households which share similar basic social and demographic characteristics are likely to be fairly similar to other households sharing those characteristics in the same part of the country. Swapping whole households reduces the possibility of creating illogical or inconsistent records and the introduction of uncertainty as to the true geographic location of high-risk households makes it more difficult for an intruder to claim with certainty to have identified a household. The method has the effect of introducing perturbations in the census data tables, but only at lower levels of geographic disaggregation. With data swapping restricted to take place within the areas defined at a certain level of census geography, any census tables produced at that level of geography or higher will be unaffected. The method is thus designed to provide disclosure control protection precisely where it is needed most, namely census tables for small areas.

Data swapping has some important practical advantages which explain its relatively widespread use. It is relatively simple to apply, particularly compared to other methods of microdata perturbation. It only has to be applied once and then all subsequent census tabulations can be produced from the perturbed microdata in the normal way. The method does, however, have some important disadvantages. It does not offer such strong protection against disclosure (for acceptable levels of perturbation) compared with post-tabular methods and this problem becomes more acute in the context of an online table generator. In addition, geographic swapping is not applicable to very small countries (see section D for further discussion of these issues).

The German Federal Statistics Office (Destatis) used a different pre-tabular method of disclosure control for their 1987 and 2011 census (the SAFE method, see Giessing and Höhne, 2010), but it has been announced that Destatis will shift to a post-tabular method for their 2021 census (Destatis, 2018). It should also be noted that for their 2020 Census, the United States Census Bureau (USCB) is moving away from data swapping as a methodology after deciding that it offered insufficient protection. The USCB are adopting the differential privacy standard for their 2020 census (see section C). The United Kingdom's Office for National Statistics (ONS) are planning to use record swapping again in their 2021 census but, in recognition of the fact that data swapping alone is insufficient (Spicer, 2020), will supplement it with a post-tabular method of disclosure control (see below).

Data swapping has actually been applied to Latin American and Caribbean census microdata which are disseminated through the Minnesota Population Center's IPUMS International project, although in this case the purpose is to protect samples of microdata rather than tabular data. Through this longstanding project, anonymized samples of census records from countries all around the world are made available on licensed release to approved researchers, including some larger Caribbean countries (Jamaica, Trinidad and Tobago and Puerto Rico). The use of data swapping to protect anonymized samples of census microdata records is a rather different application of the data swapping method, which is not directly comparable with its use for the protection of tabular data.

B. Post-tabular methods

The alternative to perturbation of the census microdata is to apply the perturbation after tabulation, that is, to perturb the data in the census tables themselves. In this way, disclosure control effectively becomes an 'add-on' to the tabulation process. There have been three main post-tabular approaches which statistical offices have employed: small cell adjustments, random rounding, and the Australian Bureau of Statistics' (ABS) cell perturbation method.

1. Small cell adjustments

As the name implies, small cell adjustments involve applying perturbations to the cells containing small counts only. Since most disclosure risk arises from small cell counts, it makes sense to apply adjustments to introduce some uncertainty about the true values of those counts. In their respective 1996 and 2001, and 1991 and 2001 censuses both the ABS and the UK's ONS applied the small cell adjustments method to protect tabular outputs. Applying small cell adjustments involves randomly adjusting small cells within tables upwards or downwards to a base using an unbiased prescribed probability scheme (Hundepool and others, 2012).

Statistical offices often release only limited information about disclosure control methods since knowledge of the methods, particularly their exact parameters, could help an intruder to 'unpick' the protection that the methods are intended to provide. However, the fact that published census tables in both countries' 2001 censuses did not include counts of 1 or 2 makes clear that this was the operative definition of 'small cells' and that cell counts taking these values were randomly perturbed to either 0 or 3, most likely with the same probabilities used for random rounding to base 3 (see below).

The small cell adjustment method is applied independently to each table. After the small cells have been adjusted, the total and marginal totals are re-calculated taking into account the adjustments made to the small cells. Therefore, the method actually perturbs small cells, totals and some marginal totals. A consequence of this is that the small cell adjustment method creates a degree of inconsistency between tables. It is relatively common that in sets of census tables which cross-tabulate different combinations of variables, there will be totals or marginal totals that appear in multiple tables. Since the small cell adjustment method is applied independently to each table, the total and marginal totals can be recalculated in different ways in each table depending on the small cell adjustments.

The removal of all cell counts of 1 and 2 provides quite strong protection against disclosure risk although the recalculation of totals in independently adjusted tables can lead to some cell adjustments being unpicked (Hundepool and others, 2012). For example, a marginal total may be adjusted in one table, but the true unadjusted value may appear in another table (if that table has no small cell counts), which can lead to the original small cell adjustment in the first table being unpicked. Since only small cell counts are adjusted the level of information loss (or reduction in data utility) is generally quite low except in highly sparse tables where the number of cell adjustments increases.

Nevertheless, both the ABS and ONS moved away from the use of this method after their 2001 censuses, although in slightly different directions. While the method removes the small cell counts themselves, it does little to prevent disclosure by differencing. Small cell counts will still be implicit as the difference between two larger cell counts. In the context of a pre-prepared set of census tables, this risk could perhaps be managed by careful review of the tables prior to publication. However, in the context of a table generator, this is likely to be impossible due to the sheer number of possible tables which can be generated.

2. Random rounding

The random rounding method, unlike small cell adjustment, is applied to cell counts irrespective of size. Counts are rounded either upwards or downwards to the nearest multiple of the rounding base, generally 3 or 5, according to prescribed probabilities. So unlike with conventional rounding, numbers are not always rounded to the nearest multiple of the base.

Random rounding is used by both Statistics Canada and Statistics New Zealand to protect their census tables. Statistics Canada uses random rounding to base 5 and have used this method since their 1971 census. Table 6 shows part of a Statistics Canada census table of the national population by age and sex, with all estimates rounded to base five. Interior cells, the total and marginal totals are all rounded independently. Original counts which are divisible by five remain unchanged while those not divisible by five are rounded either up or down according to the rounding scheme illustrated in table 7.

Table 6
Part of a Statistics Canada census table: population by age and sex, Census 2016
(Numbers of persons)

Age (in single years) and average age (127)	Sex (3)		
	Total - Sex	Male	Female
Total - Age	35,151,725	17,264,200	17,887,530
0 to 14 years	5,839,565	2,992,920	2,846,645
0 to 4 years	1,898,790	973,035	925,760
Under 1 year	369,730	189,085	180,650
1	372,615	190,900	181,710
2	378,880	193,940	184,945
3	386,200	198,615	187,590
4	391,365	200,500	190,865
5 to 9 years	2,018,130	1,034,690	983,440
5	394,530	202,500	192,035

Source: Data tables, 2016 Census, Statistics Canada (<https://www12.statcan.gc.ca/census-recensement/2016/dp-pd/dt-td/index-eng.cfm>).

Table 7
Random rounding to base five: rounding probabilities, Statistics Canada

Unit values of	Will round to count ending in 0	Will round to count ending in 5
1	4 times out of 5	1 time out of 5
2	3 times out of 5	2 times out of 5
3	2 times out of 5	3 times out of 5
4	1 time out of 5	4 times out of 5
5	Never	Always
6	1 time out of 5	4 times out of 5
7	2 times out of 5	3 times out of 5
8	3 times out of 5	2 times out of 5
9	4 times out of 5	1 time out of 5
0	Always	Never

Source: Statistics Canada (<https://www12.statcan.gc.ca/census-recensement/2011/ref/DQ-QD/conf-eng.cfm>).

Statistics New Zealand uses a very similar method in their censuses, except that rather than rounding to base five, they round to base three, an approach which they have been using since 1981. Random rounding to base 3 is applied in a very similar way. Cells counts divisible by 3 are unchanged while counts not divisible

by three are rounded either up or down using probabilities of one third and two-thirds. For example, a count of 7 would be rounded down to 6 with a probability of two-thirds and rounded up to 9 with probability one third. For a count of 8, the rounding probabilities are reversed.

Random rounding has two important advantages over conventional deterministic rounding (where numbers are always rounded to the nearest multiple of the rounding base). First, the greater uncertainty introduced by the use of probabilities, means that it is less likely that the rounding can be 'unpicked' by comparing independently rounded interior cells and totals. Secondly, the probabilities can be chosen to ensure that the sum of randomly rounded numbers is always an unbiased estimate of the sum of the original cell counts. Conventional rounding, on the other hand, will tend to create a downward bias due to the preponderance of small last digits in census data (Zeisset, 1978).

A notable consequence of the random rounding method is that tables are not strictly additive. This is clear from the first row of table 6 where the national population of Canada (35,151,725) is not strictly equal to the sum of the number of men (17,264,200) and the number of women (17,887,530), differing by five. As mentioned above, although controlled rounding methods producing additive tables do exist, the algorithms are not suited to census tables and certainly not online processing. Statistical offices applying random rounding have therefore had to educate users on the need to accept non-additive census tables as a price worth paying to enable safe publication of more detailed census outputs.

A second consequence of random rounding can be the introduction of inconsistencies between tables. If random rounding is applied independently to each table, when the same table cell appears in two or more tables, it might be rounded up in one table and down in another. For example, among the data tables from which table 6 was taken, the number of females aged 26 in Canada appears as 232,540 in some tables and 232,545 in others. However, random rounding can now be applied in a way which is consistent across tables and in their 2018 census Statistics New Zealand did indeed introduce what they called fixed random rounding. This means that where the same cell (containing exactly the same microdata records) appears in several tables, it will always be rounded in the same way. This consistency across tables is achieved using cell keys (see below).

3. ABS cell perturbation (or the cell key method)

In the publication of results from the Australian census of 2006, the Australian Bureau of Statistics developed an alternative method of perturbing cell counts. Rather than perturbing cell counts using random rounding, perturbation is applied in the form of additive noise with a certain probability distribution. Cell counts are typically perturbed by small values, say 0, ± 1 , ± 2 , ± 3 . As with random rounding, all cell counts including interior cells, the total and marginal totals are perturbed to protect against disclosure by differencing. The two methods are quite similar and random rounding can be analysed as a variant of cell perturbation.

Leaver (2009) described the properties of this additive noise as integer-value perturbations whose distribution satisfies the following criteria:

- (i) the mean is zero;
- (ii) the perturbations will not create negative cell values or very small positive cell values;
- (iii) the perturbations have a fixed variance;
- (iv) the absolute value of any perturbation is less than a fixed positive integer.

In the ABS implementation of this method, the 'very small positive cell values' referred to in the second criteria are again 1 and 2. Therefore the cell perturbation method will perturb all values of 1 or 2 to either 0 or 3 (or possibly 4 depending on the parameters) and the resultant table, like random rounding to base 3 will contain no values of 1 or 2. However, unlike random rounding to base 3, perturbed tables can contain any other whole number value other than 1 or 2.

From the outset, the method was designed to be used in the context of a flexible table generator, ABS's TableBuilder, which was launched in 2009 with data from the 2006 census. An important innovation which enabled the use of cell perturbation in a table generator was the use of cell keys. These cell keys enable cell perturbation (or random rounding) to be applied in an online table generator. Most importantly, the use of the cell key method prevents the possibility that the user of a table generator could simply request the same tabulation repeatedly, with counts perturbed in different ways each time, thereby enabling the user to observe the probability distribution of perturbed (or randomly rounded) counts and infer from this what the original counts were. The perturbation is 'fixed' so that every time a particular cell count is produced, meaning that exact combination of microdata records, the same random perturbation is applied. In this way, every time a particular table is requested, it is reproduced with all cells having exactly the same perturbed cell values every time. A further benefit of the method is that every time the same cell, for example all economically active women aged 45-49 in region 3, appears in any table, it is always perturbed in the same way. In this way, the method guarantees consistency between tables.

The method works by assigning record keys, essentially a unique code, to every record in the microdata (every person, household, dwelling or building). When a table is produced, in parallel with calculating either a weighted or unweighted count for each of the table cells, the record keys are summed for each cell to produce the cell key. This cell key is then used to seed a pseudo-random number which is used to determine the integer-value perturbation, which is applied to the cell, via a fixed lookup table (Leaver 2009). The lookup table, therefore, has a fundamentally important role to play because it determines the characteristics of the additive noise that is applied to the original cell counts by defining its probability distribution (which must satisfy the four criteria mentioned above).

Table 8 shows an example of a cross-tabulation of census data to which the ABS cell perturbation method has been applied. Two features of the table to note are that there are no cell counts of either 1 or 2 in the table although there are cell counts of 0 and 3. The original table may have contained cell counts of either 1 or 2, but the cell perturbation method would have changed those counts to take either the value 0 or 3. The second important feature to note is that, like the Canadian census data, the table is no longer strictly additive and there are discrepancies between the sums of the interior cells and the corresponding marginal totals and between the interior cells and the grand total (and between marginal totals and the grand total).

As originally implemented, the method was a two-stage process. At the first stage, perturbation was added to all table cells, including the independent perturbation of interior table cells, marginal totals and the grand total. This perturbation process produced a non-additive but protected table. The second stage of the method involved the application of an algorithm which applied a further set of adjustments to make the tables additive again. This had the effect of creating inconsistencies between tables because the algorithm had to be applied separately to each table. Where the same cell appeared in more than one table, although stage one of the method could guarantee consistent perturbation, there was no way to ensure that the additivity adjustments would be applied consistently. From 2017 onwards, the ABS changed their approach and stopped using the additivity adjustments so that now only the first stage of the method, perturbation, is applied. This meant that whereas between 2009 and June 2017, Australian census data tables were additive but with inconsistencies between tables, since June 2017 tables have been non-additive but consistent. It seems that the factors in this decision were a reassessment of the importance of consistency between tables versus additivity and also the complexity of the algorithm that re-established additivity, which can work well for smaller tables but becomes computationally very expensive, and therefore unsuitable for online processing, for larger multi-dimensional tables. The issue of non-additivity of tables is discussed further below.

More recently the ABS TableBuilder tool received some media coverage in Australia after research, first published in 2019 by Asghar and Kaafar (2020), described algorithms which they argued could unpick the cell perturbation method implemented in the TableBuilder website. The first step involved an algorithm to uncover the maximum magnitude of perturbations applied. This information was then used in a second algorithm which involved the creation of large numbers of slightly different queries from which many noisy

versions of the same single unknown original count could be calculated. This could be achieved through summation of many different breakdowns of the same total. Although the data is protected by additive noise which is cumulated in any summation, if a sufficiently large number of noisy versions of the same single original value can be calculated, then removal of the noise is possible. The attack was not actually carried out, but rather was intended to demonstrate a weakness in the TableBuilder confidentiality protections. The paper concluded by recommending differential privacy as the way to protect against such attacks (see chapter II section C on differential privacy).

Table 8
ABS cell perturbation applied to a cross-tabulation of religion by district (Belize census, 2010)
(Numbers of people)

	Corozal	Orange Walk	Belize	Cayo	Stann Creek	Toledo	All districts	Discrepancy
Anglican	382	485	9 737	1 496	1 545	246	13 894	3
Bahai Faith	26	14	44	68	45	6	200	-3
Baptist	619	421	3 856	1 850	1 764	2 400	10 912	2
Buddhism	58	104	366	162	48	7	744	-1
Hinduism	104	95	336	41	22	5	606	3
Islam	66	0	425	49	27	5	568	-4
Jehovah's Witness	710	1 006	1 432	1 081	505	335	5 071	2
Mennonite	2 536	4 947	331	2 852	206	621	11 491	-2
Methodist	525	293	6 125	369	860	577	8 747	-2
Mormon	131	189	385	500	67	7	1 277	-2
Nazarene	824	150	1 905	2 605	1 260	1 852	8 591	-5
Pentecostal	1 720	2 035	5 079	10 844	2 845	3 120	25 643	0
Rastafarian	14	18	248	113	94	35	520	-2
Roman Catholic	18 331	19 940	33 260	24 264	12 878	13 259	121 936	4
Seventh Day Adventist	3 977	2 306	5 377	3 000	1 364	465	16 491	2
Salvation Army	12	7	242	102	8	3	366	-8
Other	2 729	7 154	5 434	6 466	2 405	3 120	27 310	2
None	6 482	5 478	12 086	13 680	5 527	3 806	47 057	-2
All religions	39 243	44 651	86 666	69 539	31 463	29 865	301 431	4
Discrepancy	-3	9	-2	-3	-7	-4	7	-6

Source: ECLAC on the basis of data provided by the Statistical Institute of Belize

Note: Unweighted data.

The ABS responded, reporting that “since early 2017, the ABS has been working with Dr Dali Kaafar and his co-researchers [and] as part of their research, they constructed a very elaborate and obscure attack scenario. The researchers brought this theoretical risk to the attention of the ABS and we have worked with them to address it. The ABS has since then put in place measures to mitigate this potential vulnerability suggested by the researchers” (ABS, 2019). This involves monitoring usage of the TableBuilder website to detect suspicious series of queries characteristic of an attack. ABS also produced their own research arguing that the likelihood of a breach of confidentiality was very low (Bailie and Chien, 2019).

This serves to emphasise that simulation attacks are an important part of the process of development of disclosure control and privacy protection methods and that risk cannot be completely eliminated. The fact that this research demonstrated a theoretical risk of disclosure should not be taken as a sign that TableBuilder is not, in ABS’s words “a world-leading statistical tool.” Statistical offices which are at the forefront of the development of statistical disclosure control methods, and actively involved in research in this area, are by virtue of that engagement with experts in data privacy likely to subject their own methods to critique and scrutiny as part of the research and development process. This is also borne out by the fact that, since its development and implementation, the ABS’s cell key perturbation method has attracted interest from other statistical offices who are increasingly seeking to provide their census users with online table generators, and so are looking for an appropriate disclosure control solution which is applicable in this context.

In preparation for the 2020 census round, two EU-funded research projects (Eurostat, 2020) sought to identify best practices for disclosure control of census data and to provide open source tools for the application of selected perturbative methods. Two methods of disclosure control were recommended: the pre-tabular method of targeted record swapping and the post-tabular random noise method (ABS cell perturbation). With the two methods having different pros and cons, this allows countries to use either or indeed both methods in accordance with national circumstances and preferences. In applying the methods, statistical offices will also have flexibility to select the parameters which determine, for example, the extent of data perturbation, which is likely to vary depending on national data (and also whether countries are applying one or both disclosure control methods). Tools were developed to facilitate the testing and ultimately implementation of these methods, including SAS code for each of the respective methods and there has also been significant progress towards implementation of these methods in specialist SDC software (μ -Argus, τ -Argus, sdcMicro, sdcTable).

Based on what some statistical offices have indicated about how they intend to approach the issue of statistical disclosure control in their upcoming censuses, it appears there will indeed be some variation in the way that different countries choose to apply these recommendations. For example, the German Federal Statistical Office (Destatis) have announced their intention to apply the ABS cell perturbation method in their 2021 census, stating that "in view of the advantages, i.e. the high accuracy of the procedure and the cross-tabular consistency of cells with identical content, non-additivity seems not much more than a cosmetic flaw" (Destatis, 2018). The cell key method will replace a pre-tabular perturbation method which Destatis used in their 2011 census. For Destatis, a major advantage of the cell perturbation method over the previous pre-tabular method is the lower level of data perturbation involved and therefore reduced information loss (Giessing and Höhne, 2011).

The UK's ONS will also use a variant of the ABS cell perturbation method in their 2021 census. However, unlike Destatis, ONS will continue to use targeted record swapping, an approach that was used in the UK censuses of both 2001 and 2011. The planned use of the cell key perturbation method is described as 'light touch,' reflecting the fact that it will be used in conjunction with targeted record swapping (Spicer, 2020; Blanchard, 2019; and Dove, Blanchard and Spicer, 2017). This means that cell perturbation will be applied to a relatively small proportion of cells which will be perturbed by no more than ± 1 and census tables will contain cell counts of 1 and 2 (albeit having introduced uncertainty about whether these are the true original counts).

The rationale for this approach is that while record swapping does not provide adequate protection on its own, used in combination with 'light touch' cell perturbation, the two methods combined provide sufficient protection. Targeted record swapping, it is argued, introduces uncertainty to small cell counts, while cell perturbation protects against disclosure by differencing (Spicer, 2020). This 'light touch' cell perturbation makes it possible to retain the additivity of tables because the total and marginal totals can be recalculated simply by summing the perturbed interior cells of the tables. With a more normal level of cell perturbation, the total and marginal totals cannot be re-calculated in this way, because the loss of accuracy in the total and marginal totals would be too great. The total and marginal totals are effectively being perturbed by the sum of all the perturbations applied to their corresponding interior cells which is only viable with very 'light touch' perturbation. However, as with the ABS algorithm for restoring additivity, this creates the problem of inconsistency between tables because the re-calculation of the total and marginal totals will be different in each table depending on the cell perturbations in those tables.

Further facilitating the testing and use of the cell key perturbation approach are the recently developed R packages cellKey (Meindl, 2020; and Meindl and Enderle, 2019) and ptable (Enderle, 2019) which implement a refined version of the method. The cellKey package generalises the ABS cell perturbation method, including a new specification for record keys, while ptable is a tool for calculating the lookup tables which the cellKey package then uses to generate and apply additive noise to tables. The cell key method has also been implemented in τ -Argus, a software package mainly used by SDC experts partly developed with EU funding.

C. Differential privacy

Differential privacy is a relatively new approach to the protection of statistical databases. It was originally developed by experts in cryptography, which is the use of mathematical methods of encryption to protect information. Differential privacy itself is really a mathematical definition of privacy formulated in terms of the information that can be learned (or not) about any single individual in a database. Different methods of protecting data can be designed to meet this standard. These algorithms protect privacy through data perturbation and therefore there is no escaping the risk-utility trade-off, which is just as fundamental to differential privacy as it is to traditional statistical disclosure control. However, the SDC approach to measuring and managing risk depends quite heavily on assumptions and judgements about the ways in which a breach of confidentiality could occur. For example, where traditional statistical disclosure control makes assumptions about the information available to an attacker, differential privacy makes no such assumptions. Statistical disclosure control is concerned primarily with the most plausible threats to privacy, while differential privacy accounts for the worst case scenario. A consequence of this is that differential privacy will tend to lead to a rather different trade-off between risk and utility: a stronger guarantee of privacy but at a somewhat greater cost to data utility.

Differential privacy provides a mathematical guarantee that anyone viewing an output from a database will 'essentially' make the same inference about any individual's private information, whether or not that individual's private information was included in the calculation of the output. What can be learned about an individual as a result of their private information being included in a differentially private analysis is limited and quantified by a privacy loss parameter, usually denoted ϵ (Wood and others, 2018). When $\epsilon = 0$, this equates to a provable guarantee of the complete impossibility of learning anything about an individual in the database, but unfortunately this also means that the outputs returned to the user are independent of the data and consist purely of noise. In practice, ϵ is set to a small positive value. A low value of ϵ equates to a strong guarantee of privacy but with greater loss of accuracy. Increasing the value of ϵ slightly improves data utility but with a somewhat higher level of disclosure risk. An additional advantage of differential privacy is that the method and parameters used to confidentialise outputs can be made public without undermining the privacy guarantee. This facilitates a wider understanding of the trade-offs involved in privacy protection and enables analysts to account, to some extent, for data perturbation in their analyses.

Differentially private methods tend to be mathematically and/or computationally demanding and have yet to be widely applied in official statistics. Most existing mechanisms were created for computer science applications, not the needs of official statistical agencies. For example, there is currently no accepted mechanism for applying differential privacy to the results of stratified probability samples of the kind used in official household surveys (Garfinkel, Abowd and Powazek, 2018). Expertise in differential privacy is currently concentrated primarily in academia, with smaller, but growing, communities in government and industry (Page, Cabot and Nissim, 2018). However, this could be set to change since the US Census Bureau has confirmed that their 2020 Census results will be protected using differential privacy, which they describe as "the new gold standard in data privacy protection." The decision to adopt differential privacy reflected a concern that the record swapping approach applied in the previous three censuses and other methods of statistical disclosure control did not provide sufficient protection of privacy in the face of the threat posed by modern computing power and the potential for linkage to external databases.

Following the 2010 Census, the Bureau conducted an internal experiment to see if it was possible to, first, reconstruct the 2010 microdata records from the tabular data that had been released and, second, to link the reconstructed records to a commercially available database, in order to assess how feasible it was to identify individuals. In this experiment, individual records for census block, sex, age, race and ethnicity were successfully reconstructed for 46 per cent of the population; and block, sex, and age were then linked to commercial data, which provided putative re-identification of 45 per cent of the population. Subsequent verification against the original census data confirmed 38 per cent of these matches were correct, equating to

17 per cent of the population. For comparison, the previous time that a similar experiment had been carried out (on the American Community Survey), the putative re-identification rate was 0.017 per cent, and the proportion of those confirmed correct, 22 per cent (Leclerc, 2019). It was based on these results, that it was decided that the more robust privacy standard was needed for Census 2020. This will include provable bounds on the accuracy of the best possible database reconstruction given the released tabulations (Garfinkel, 2017).

Fundamental to the Census Bureau's implementation of differential privacy in Census 2020 is the fact that the United States uses a short census form with the majority of 'census' variables now collected through the American Community Survey, a continuous sample survey. The short census form only collects five tabulated variables (age, sex, race, ethnicity and relationship to householder) plus location. So, despite its size, the decennial census is actually the easiest US Census Bureau product to make differentially private (Garfinkel, 2019).

The algorithms that will be used to implement differential privacy are in the latter stages of development and testing. They do not fall neatly into one of the two categories used in this paper to distinguish between methods of statistical disclosure control: pre-tabular methods and post-tabular methods. In the first instance, noise is applied to frequency counts. But then using a top-down algorithm, which works from national-level down to state, county, tract, block group and finally block-level, perturbed counts are allocated back to geographic areas in a way which allows for the creation of a new differentially private microdata file: the microdata detail file (MDF). All census tabulations are then produced from this MDF. Records in the MDF have no exact counterpart in the confidential data. The algorithms preserve some very basic counts such as the number of people per block which remain unperturbed as required by law (USCB, 2011). All other counts are potentially subject to perturbation, with the algorithms applying appropriate levels of perturbation at different levels of geography (more perturbation for finer geographies).

D. Applicability of SDC methods to Caribbean censuses

Three approaches to disclosure control have been considered: pre-tabular methods, post-tabular methods and differential privacy. Differential privacy offers a rigorous privacy standard, is contributing much to our understanding of the disclosure control problem and may become an important part of official statistics in the future. However, at the time of writing, no statistical office has published differentially private census results. The US Census Bureau will publish their 2020 census results using differential privacy, but this is for a short form census with complete enumeration of five tabulated variables. Even for these variables, the methods being used are complex and computationally intensive. No statistical offices that use a long census form (collecting, say, 100 or 200 variables) are close to adopting differential privacy at the present time and it is not a feasible solution for Caribbean censuses.

Pre-tabular methods do have some clear practical advantages which is why at least some statistical offices have preferred this approach in recent census rounds. The principal advantage is that perturbation is applied once to the microdata and then most or all subsequent census outputs can be produced in a consistent way from the perturbed microdata. Pre-tabular methods do not create any problems of either non-additivity or consistency between tables.

There are, however, a number of disadvantages to the pre-tabular methods. Post-tabular methods provide more control on the perturbation of small cells compared to pre-tabular methods (Shlomo and Young, 2008). Pre-tabular methods do not generally remove small cell counts in the way that the post-tabular methods can. This means that intruders may still be able to claim that they have identified individuals from these small counts. The data producer must then rely on the argument that, having created uncertainty about whether the small cell counts are actually the true cell counts, the intruder cannot be sure that their identification or disclosure of information is in fact correct. This is a less strong level of protection, than that provided by the post-tabular methods which can actually remove small cell counts (and the ability to calculate them indirectly).

Rinott and others (2018) commented on the increased risk of inferential disclosure associated with flexible table generators in which tables can be manipulated and differenced and how this typically called for perturbative methods such as rounding or additive noise applied to the cell counts. In their review of disclosure risk and data utility in flexible table generators, Shlomo, Antal, and Elliot (2013) concluded:

"Whilst statistical agencies can argue that there is a level of uncertainty in the table arising from pre-tabular methods such as record swapping, for a flexible table generating server of census data made freely available on the internet, intruders can easily re-identify the characteristics of the unique individuals. In addition, these methods made little impact on the disclosure risk of differencing tables. Applying preliminary rules and an SDC method on the final output table provided the lowest disclosure risk with the highest utility...to rely solely on SDC methods implemented 'on the fly' for protecting user-defined tables means that some form of random noise needs to be applied to the output tables."

It should also be noted that record swapping, the simplest and most commonly used pre-tabular method, is not applicable to small countries. By definition, data swapping provides protection only below a certain level of geographic disaggregation defined by the areas within which records are swapped (swapping was within states in recent US censuses), but it provides no protection at that level or above where swapping has no effect. In larger countries, this means that protection is applied where it is needed: for more geographically disaggregated statistics. However, for Caribbean countries there is no distinction that can be drawn between small area statistics requiring protection and higher-level geographies where the tables are safe. For most Caribbean countries, even national level tabulations would be regarded as small area statistics requiring disclosure control protection. Therefore, a method which provides protection at all levels of geographic disaggregation is required.

A further argument against the use of pre-tabular methods for Caribbean censuses, is that the application of such methodologies is a relatively complex undertaking, that requires expertise in statistical disclosure control. Even if a common approach were developed, it would have to be tailored to each individual country. It would add a complicated additional stage to census data processing which, bearing in mind the challenges experienced in publishing and disseminating outputs in previous census rounds, would not be practical. It would also add an element of risk to census processing (the paper "When Excessive Perturbation Goes Wrong and Why" by Cleveland and others, 2012, offers a cautionary tale).

By contrast, the post-tabular methods lend themselves more readily to automation. If an automated routine is available, the statistical office has to do nothing more than select exactly which variant of the method they want to use, select the appropriate parameter(s), and apply the method. The methods can be applied in the same way to any census, from the census of Montserrat to the census of Jamaica (or indeed the census of Brazil). With the use of the REDATAM software for tabulation and online dissemination well-established in the region, post-tabular methods of disclosure control could be implemented as part of the software's tabulation functionality. In this way, the methods could be made available for use by the countries of the region both for the production of their regular published census outputs and for use in the online applications built using REDATAM.

The obvious disadvantage of post-tabular methods is the loss of additivity of census tables, which probably explains why they have not been more widely adopted. Unless they are familiar with the census data published by Canada, New Zealand or Australia, both producers and users of census statistics are accustomed to the fact that tables are additive. However, the relatively long-standing use of random rounding by the statistical offices of Canada and New Zealand and the more recent adoption of cell perturbation by the ABS and now the German Federal Statistical Office suggests that it is possible to overcome whatever initial adverse reaction there may be to the use of these methods. For a small loss of accuracy and additivity, post-tabular perturbation allows statistical offices to safely publish much more detailed and disaggregated census outputs than is otherwise possible. If this is properly explained to users of census data, then there is no reason to think they would not be just as prepared to accept this as census data users in the four countries mentioned above.

As mentioned, the United Kingdom is pursuing a different approach and a desire to maintain additivity of tables appears to be one of the main reasons why. As mentioned previously, the ONS will continue to apply record swapping combined with a 'light touch' application of the cell perturbation method. This dual application of both pre- and post-tabular methods does enable additivity to be restored by summing marginal totals and totals from the lightly perturbed interior cells, but at the expense of consistency between tables. It is not clear that additivity is preferable to consistency. 'Fixed' random rounding or cell perturbation with cell keys might not produce additive tables but at least every census statistic has a single published value. Insisting on additivity leads to publishing multiple different values for the same statistic which, for users, begs the question which of the multiple values they should use.

III. Comparative analysis and testing of methods

This chapter analyses in more detail the post-tabular disclosure control methods, random rounding and cell perturbation, which are the most appropriate and feasible for application to Caribbean censuses. It also reports on the testing of these methods that was carried out, as part of this study, using census data for Jamaica, Belize and Saint Kitts and Nevis.

A. Comparative analysis of methods

Fundamental to the cell perturbation method is the transition matrix which defines the probability distribution of the random perturbations which are applied to the cell counts. Diagram 1 illustrates two possible transition matrices. The transition matrix determines the probabilities with which an original frequency i is replaced with a perturbed frequency j . Taking transition matrix A, it can be seen that an original frequency of zero will be preserved with probability 1. All the methods considered here have this property, that original zeros are preserved. Non-zero values may be perturbed to zero but zero values are never perturbed to become non-zero. Zeros in frequency tables can be distinguished between: non-structural zeros (for example the number of usual residents of Saint Lucia born in Ukraine) which may happen to be zero but are not automatically zero by definition; and structural zeros that are zero by definition (for example, unemployed persons aged under 15). While non-structural zeros could plausibly be perturbed, the perturbation of structural zeros would result in illogical counts. In practice, simply preserving all original zeros is the most practical option because then there is no need to distinguish between structural and non-structural zeros.

For an original count of 1, there is only one possible negative perturbation which is applied with a probability of 0.38 resulting in a perturbed count of 0. An original count of 1 is preserved with probability 0.38, is replaced by a 2 with probability 0.137, and so on. For small cell counts (in this case 1s and 2s), the probability distribution of the perturbations is asymmetrical in order to avoid introducing negative frequency counts, but for both original counts of one and two, the expected value of the perturbation is still zero so that no overall bias is introduced to the table. For original cell counts of three, a symmetric probability distribution can be used, with the original cell count being perturbed to any value between 0 and 6. For all original cell counts

greater than 3 (which are not shown in transition matrix A), the same symmetric probabilities are used so a cell count of 10, for example, would be perturbed within the range 7 to 13 inclusive.

Diagram 1 shows a second perturbation scheme (B) which differs in one important respect from scheme A. In scheme A, the parameter js is set to zero which means that all whole number counts are admissible as perturbed counts. In scheme B, $js = 2$ which means that small counts of 1 or 2 are removed entirely from the table. As can be seen from the transition matrix, counts of 1 are perturbed to 0 with probability 0.667 and perturbed to 3 with probability 0.33 (and to 4 with negligible probability). For original cell counts of 2, the probabilities are almost exactly reversed, with 0.334 of 2s being perturbed to zero and 0.663 of 2s being perturbed to 3. With this scheme eliminating all counts of 1 and 2, the noise distributions for original cell counts of between 1 and 5 are also asymmetric (but again, zero in expectation). As with scheme A, the first symmetrical probability distribution is also applied to all larger original frequencies.

Diagram 1
Two alternative transition matrices for cell perturbation

A. Transition Matrix: cell perturbation ($D=3$, $V=1.5$, $js=0$, $p_{stay}=0.4$)

		j (target frequency)						
		0	1	2	3	4	5	6
		0	1					
i (original frequency)	j	1	0.38	0.38	0.137	0.069	0.035	
i	2	0.123	0.193	0.4	0.163	0.088	0.033	
i	3	0.024	0.087	0.189	0.4	0.189	0.087	0.024



B. Transition Matrix: cell perturbation ($D=3$, $V=2$, $js=2$, $p_{stay}=0.33$)

		j (target frequency)							
		0	3	4	5	6	7	8	9
		0	1						
i (original frequency)	j	1	0.667	0.33	0.002				
i	2	0.334	0.663	0.002	0				
i	3	0.149	0.507	0.257	0.075	0.013			
i	4		0.38	0.38	0.137	0.069	0.035		
i	5		0.175	0.185	0.333	0.141	0.101	0.064	
i	6		0.044	0.107	0.183	0.333	0.183	0.107	0.044



Note: Produced using the R package ptable (Enderle, 2019).

The transition matrix A can therefore be seen as providing protection of small cells in a somewhat similar way to record-swapping. It protects small cells by creating uncertainty about the true original cell count. Transition matrix B, on the other hand, protects small cell counts by removing them altogether from the table.

These transition matrices were calculated using the R package ptable (Enderle, 2019) which was created for precisely this purpose. There are four parameters which define the transition matrix. The parameter D defines the maximum perturbation, with $D = 3$ both for perturbation schemes A and B. The parameter V defines the variance of the perturbations. This variance (or equivalently standard deviation) is the parameter that really determines the amount of noise or perturbation that is applied to the data. The parameter 'pstay' sets the value at the centre of the symmetrical probability distribution i.e. the perturbation pattern applied to all except the smallest cell counts. This parameter, which takes the value 0.4 in transition matrix A and 0.333 in transition matrix B, determines the proportion of (non-small) original cell counts for which the perturbation is zero.

Once those parameters are defined, they become constraints in an optimisation problem which is to find the probability distribution which maximises the entropy of the perturbations given the parameters D , V and $pstay$. The package ptable solves this optimization problem. Entropy refers to the 'spread' of the probability distribution and maximising this property means making the distribution as even (i.e. as close to a uniform distribution as possible), subject to the constraints D , V and $pstay$. In practice, these other parameters are generally set in such a way that the resulting probability distribution resembles a normal or Laplace distribution (in discrete form) rather than a uniform distribution.

The rationale for the entropy criteria is that a table in which interior cell counts are evenly spread across the different categories of the variables forming the table has lower disclosure risk than a table in which the counts are clustered in one or a small number of cross-classified categories (in which small cell counts and zeros are more likely to be prevalent). Cell perturbations should therefore contribute to increasing rather than reducing entropy.

Rinott and others (2018) note the intuitive appeal of using maximum entropy subject to variance to determine the pattern of perturbation although they point out that considering the problem through the lens of differential privacy suggests that entropy might not be the best criteria to use. In the field of differential privacy, the Laplace distribution is commonly used to generate noise and their results suggest that discrete Laplace noise provides more efficient protection of statistical confidentiality than the discretised version of normally distributed noise (which has approximately maximum entropy). Nevertheless, the three other parameters D , V and $pstay$ still provide very significant scope for creating probability distributions of quite different forms and the maximum entropy criterion is only one of four elements which together determine the probability distribution.

The random rounding methods described above are essentially quite similar to cell perturbation and can be analysed as a variation of the cell perturbation method. There is an implied probability distribution for the rounding adjustments⁴ and that implied distribution can be used to compare random rounding (base 3 or 5) with various parameterizations of the cell perturbation method.

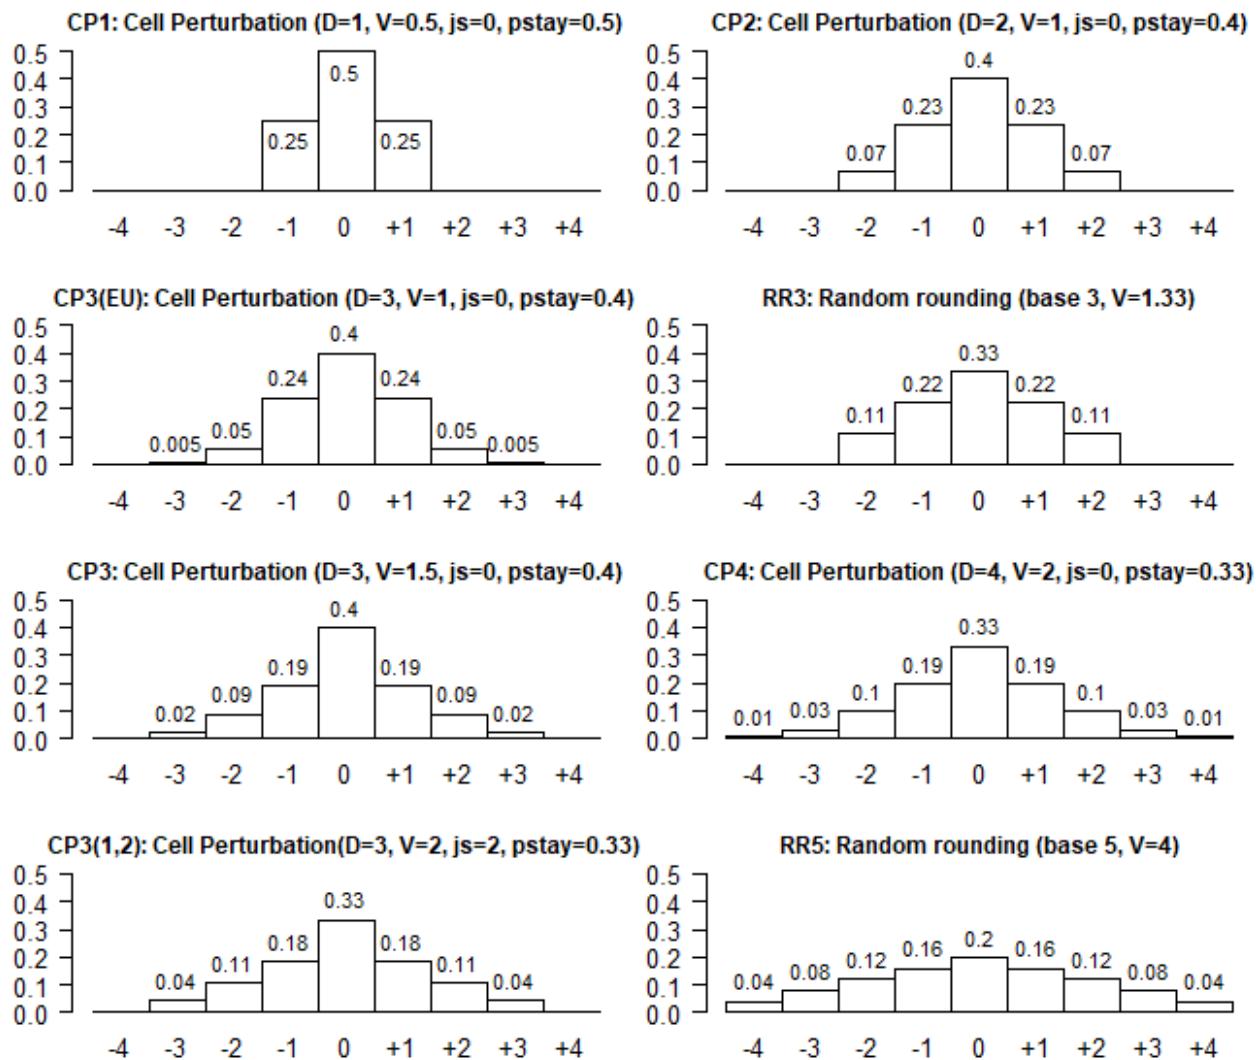
Figure 1 illustrates the symmetric probability distribution for eight different perturbation or rounding schemes which were considered in this comparative analysis, that is, the perturbations which are applied to all cell counts other than the very smallest. The perturbation schemes are ordered from left to right and then from top to bottom in order of increasing variance of the perturbations. Scheme CP1 (cell perturbation with maximum magnitude $D = 1$) involves the 'lightest' perturbation of the data with the added noise having the lowest variance V . Scheme RR5 (random rounding to base 5) is the 'heaviest' perturbation of the data in which the noise has the highest variance.

The variance is one of the parameters used to specify the cell perturbation schemes although it is more useful to think of the standard deviations, rather than the variances, as being indicative of the amount of additive

⁴ This refers to the distribution of rounding adjustments assuming all cell count remainders on division by the rounding base are equally likely. This will be approximately true except for among small cell counts and so these implied distributions of rounding adjustments can be compared to the symmetric distributions of cell perturbations which are applied to non-small cells.

noise that is being applied. For example, random rounding to base 3 implies perturbations having a variance of 1.33 and random rounding to base 5 implies perturbations having a variance of 4 but it is the respective standard deviations of 1.15 and 2 which provide a better idea of the amount of additive noise that is being added to the data i.e. random rounding to base 5 applies about 75 per cent more perturbation than random rounding to base 3.

Figure 1
Post-tabular perturbation schemes for census data (the symmetric case)



Note: The random rounding schemes are specified in terms of their rounding probabilities rather than the parameters D, V, js, and pstay. However the implied parameter values for random rounding to base 3 would be $D = 2, V = 1.33, js = 2$ and $pstay = 0.33$, and for random rounding to base 5 would be $D = 4, V = 4, js = 4$ and $pstay = 0.2$. Produced using the R package ptable (Enderle, 2019)

To create perturbation schemes which apply increasing levels of perturbation, the parameters D and V are increased. If the intention is to increase the level of perturbation, the parameter pstay (the proportion of zero perturbations) is reduced. For the perturbation schemes in figure 1, therefore, the values of D and V gradually increase from scheme CP1 up to scheme RR5 while the values of pstay generally decrease. For the cell perturbation schemes, js = 0 except for CP3(1,2) for which js = 2. So only schemes CP3(1,2) (and the random rounding schemes) would see the removal of all cell counts of 1 or 2. The parameter setting

$js = 2$ implies, for example, cell counts of 1 and 2 are perturbed to either 0 or 3 or 4. Cell counts of 3 cannot be perturbed down to 1 or 2 so must be perturbed down to 0. The parameter setting of $js = 2$ is incompatible with parameter settings of $D < 3$ and $V < 2$ because a certain minimum level of perturbation is required in order to be able to remove the cell counts of 1 or 2.

Scheme CP1 applies a minimal level of perturbation, insufficient to serve as a primary method of disclosure control. The other schemes broadly reflect the range of different levels of perturbation which have been tested or applied by statistical offices and which could be appropriate for Caribbean census. Scheme CP3(EU) is that recommended in the EU funded project on harmonised disclosure control for European censuses. Scheme RR3 is randomised rounding to base 3 as used by Statistics New Zealand. Scheme RR5 is randomised rounding to base 5 as employed by Statistics Canada. CP3(1,2) shares with the ABS method the property that all counts of 1 and 2 are suppressed.

B. Testing of selected post-tabular methods

1. Measurement of disclosure risk and data utility

The fundamental goal of statistical disclosure control is to reduce disclosure risk to an acceptable level while minimising the information loss that is necessary to achieve that. The two criteria against which methods of disclosure control are assessed are therefore disclosure risk and data utility. Measuring the impact of disclosure control methods for tabular data on data utility is relatively straightforward. The greater the perturbation that is applied, the greater the damage to the data and the lower its utility. Measures of data utility therefore measure the distance between the original data and the perturbed or masked data. In line with other similar studies (see Antal, Enderle and Giessing, 2017 and Shlomo and Young, 2005), the following three measures of data utility have been used: the average absolute distance, the average relative distance and the Hellinger's distance. Formal definitions of these measures are provided in the technical annex. In summary, the average absolute distance is an additive measure of the average distance between the counts in the original and the perturbed table. The average relative distance is the equivalent proportional (or multiplicative) measure. Hellinger's distance is a measure of the difference between the square roots of the counts in the original and perturbed table. The Hellinger's distance takes some account of both the distance between original and perturbed counts and their size. It gives greater weight to a given distance between two small cell counts than the same distance between two large cell counts and unlike the relative absolute distance, it is defined for original counts of zero (Shlomo, 2018).

Whereas assessing data utility for tabular data is relatively straightforward, measuring disclosure risk is more challenging. Unlike with microdata for which there are fairly well-established methods for analysing disclosure risk, there is less consensus on how disclosure risks for tabular data should be evaluated and quantified and various methods have been proposed (see Chipperfield, Gow and Loong, 2016). The simplest approach to assessing disclosure risk is to consider low counts to be problematic and to set some threshold below which unperturbed low counts are deemed to be problematic (Duncan, Elliot, and Salazar-Gonzalez, 2011). Notably this doesn't account for the way that small cell counts can be deduced by differencing larger cell counts, particularly in the context of a flexible table generator where the quantity of tables which are potentially available to users makes it more important to protect against disclosure by differencing.

Two measures of disclosure risk were used to test the various post-tabular methods. The first assesses the different methods of post tabular data perturbation based on the success with which they mask the true values of the counts in the original table. The random perturbation of counts introduces uncertainty about the true original value for all cells. However, if an intruder has at least some knowledge of the perturbation scheme that has been applied, using this information together with the information contained in the additive relationships between the original interior cell counts and original marginal totals, it is possible to deduce, at least, upper and lower bounds for the true cell count and sometimes the true value itself can be deduced. This is clearly an undesirable characteristic of methods seeking to mask true cell counts.

Following the same approach as Giessing (2016) and Enderle, Giessing and Tent (2018), a modified version of the shuttle algorithm originally proposed by Buzzigoli and Giusti (1999) was used to calculate upper and lower bounds, that is, a feasibility interval for the original cell counts. The intruder is assumed to have a perturbed table, together with some knowledge of the disclosure protection method which has been applied (specifically the maximum perturbation applied to any single count, for example ± 3), while the original counts are all assumed to be unknown. The perturbed table and the maximum perturbation immediately provide initial estimates of upper and lower bounds for all cell counts. The shuttle algorithm then uses the additive relationships between interior and margin cells and the implications that these have for the respective upper and lower bounds of each cell, to recursively refine these initial estimates. The algorithm converges after a few steps with final estimates of upper and lower bounds for each interior and total cell which may be closer together than the initial estimates. If the final upper and lower bound are equal, the original cell count is disclosed.

The second measure of disclosure risk is not based on such a clearly defined threat but is a composite indicator of risk which combines several separate indicators each intended to capture different aspects of disclosure risk. This measure of disclosure risk is a simplified version of that proposed by Shlomo, Antal and Elliot (2015). It is based on a weighting of three separate indicators designed to measure different characteristics of census tables which are associated with disclosure risk in the following ways: (a) small cell values have higher disclosure risk than large values; (b) uniformly distributed frequencies imply low disclosure risk; (c) the more zero cells in the census table, the higher the disclosure risk.

Unlike the first measure of disclosure risk described above, this measure of risk is calculated for the original table and the perturbed table (obviously the risk should be lower for the perturbed table than the original table). The measure is constructed in such a way that it is bounded by 0 (low risk) and 1 (high risk) and is denoted $R(F, w_1, w_2)$. Further details of both risk measures are provided in Annex 2.

2. A description of the test data

In order to further evaluate the applicability of the disclosure control methods described in the previous chapter to Caribbean census data and, potentially, census data from other countries in the region, comparative testing was carried out on data from three countries: Jamaica, Belize and Saint Kitts and Nevis. These three countries were chosen in order to test the methods with data from one large Caribbean country, one medium-sized Caribbean country and one small Caribbean county. The test data was designed to be representative of data tables which could be made available through REDATAM applications, subject to disclosure control considerations. For each of the countries, two-way, three-way, four-way and five-way cross-tabulations were created, as could be made available through REDATAM's cross-tabulation functionality which permits the cross-tabulation of up to five variables.

National census reports publish such cross-tabulations although less frequently four- or five-way cross-tabulations. These four- and five-way cross-tabulations can also be thought of as hypercubes which are large multi-dimensional cross-tabulations, often too large to be conveniently displayed as a traditional cross-tabulation, but from which analysts might select a more manageable subset of data according to their needs.

It should be remembered that for the purposes of disclosure control, the number of dimensions of a cross-tabulation is not defined just by the variables which span the rows and columns of the table, but also any variables which determine the universe or total population for the table. For example, a cross-tabulation of occupation by level of education for women would look like a two-way cross-tabulation of occupation by education but since the sex variable defines the population for the table, it should be regarded as a three-way cross-tabulation, or at least part of the full three-way cross-tabulation of occupation by education by sex. It is very common that analysts require three, four and five census variables to fully specify the data that they need, either to define the population or to cross-classify that population. For example, an analysis of the impact of non-communicable diseases on labour market participation requires, at least, one health variable and one labour market variable. The analyst is also very likely to want to know how age and sex bear on this issue, and perhaps geography or household structure, therefore very quickly the analyst requires at least four

cross-classifying variables. The three-, four- and five-way cross-tabulations created here should be understood as representing this need that analysts have for multi-dimensional cross-tabulations, whether or not they necessarily need every cell in the hypercube (i.e. every possible combination of categories of the cross-classifying variables).

Table 9
Cross-tabulations of census data and disclosure risk

	Two-way: Age x Sex	Three-way: Age x Sex x Ethnicity	Four-way: Age x Sex x Ethnicity x Place of residence	Five-way: Age x Sex x Ethnicity x Place of residence x Place of birth
Number of interior cells in the cross-tabulation (K)				
Jamaica 2011	194	1 164	16 296	228 144
Belize 2010	194	2 910	49 470	346 290
Saint Kitts and Nevis 2011	194	1 940	27 160	162 960
Average cell count (N/K) ^a				
Jamaica 2011	11 414	1 878	134	9.43
Belize 2010	1 563	104	6.12	0.87
Saint Kitts and Nevis 2011	243	24	1.74	0.29
Percentage of cell counts of 1 or 2				
Jamaica 2011	0	3.3	20.5	10.0
Belize 2010	0	11.4	19.5	8.1
Saint Kitts and Nevis 2011	0	17.1	8.9	3.1
Percentage of cell counts of 0, 1 or 2				
Jamaica 2011	0	5.1	54.1	85.5
Belize 2010	0	21.4	76.9	95.6
Saint Kitts and Nevis 2011	0	62.5	90.0	98.1
Disclosure risk R(F, w ₁ , w ₂) ^b				
Jamaica 2011	0.01	0.06	0.12	0.2
Belize 2010	0.01	0.08	0.15	0.23
Saint Kitts and Nevis 2011	0.01	0.17	0.24	0.3

Source: ECLAC on the basis of population and housing census data provided by the Statistical Institute of Jamaica, the Statistical Institute of Belize and the Department of Statistics, Ministry of Sustainable Development, Saint Kitts and Nevis.

^aN is the population total for the table equal to the sum of the interior cell counts.

^bSee Annex 2 for a definition of this measure of disclosure risk.

Tables of test data were created for each of the three countries using the same five variables: age, sex, ethnicity, place of residence and place of birth. To illustrate the effect of country size, table 9 shows some basic characteristics of the test data and indicators of disclosure risk.

A number of observations can be made of these cross-tabulations. As more variables are cross-tabulated, the number of cells increases rapidly (the total number of interior cells being the product of the number of categories in each cross-classifying variable). This implies that the average cell size decreases rapidly. A two-way cross-tabulation of single year of age by sex (with some minor top coding of age) produces an (unweighted) average cell count of 11,414 for Jamaica, 1,563 for Belize and 243 for Saint Kitts and Nevis. For a five-way cross-tabulation of all the variables, these figures become 9.4, 0.87 and 0.29 respectively. As the average cell count decreases, there are more zeros and small cell counts in the table and therefore disclosure risk increases. This is reflected in the disclosure risk measure $R(F, w_1, w_2)$ which increases with each successive cross-tabulation (or disaggregation) of the data.

The results in table 9 illustrate how quickly small cell counts become apparent and tables become sparse with more detailed disaggregations of the data. These disclosure risks appear somewhat more quickly in data for a small country such as Saint Kitts and Nevis. However, this does not mean that there are no risks for larger countries. In the three-way cross-tabulation of age, sex and ethnicity for the Jamaica census dataset, more than three per cent of cell counts are either 1 or 2 which is less than for Belize and Saint Kitts and Nevis, but that number of small unprotected cell counts could still prevent the table being published.

Table 9 also makes clear how perturbative methods of disclosure control make it possible to publish much more detailed and disaggregated census tables. In the absence of such methods, even the three-way cross-tabulations in the test data may not be published because of the presence of small cell counts, depending on the operating practice of the statistical office concerned. However, the introduction of a small amount of perturbation to mask those small cell counts (and the small cell counts which are implicit through differencing) could enable release of all the three- and four-way cross-tabulations and the five-way cross-tabulation for Jamaica. The five-way tables for Belize and Saint Kitts and Nevis are probably too sparse for publication (see chapter IV, section B).

Since population size is such a fundamental determinant of disclosure risk when it comes to census tables, the implications for larger countries are perhaps best understood by thinking in terms of subnational populations of a similar size to the Caribbean countries discussed here. To take one example, release of census data for the metropolitan area of Cali in Colombia is likely to create broadly similar disclosure risks and therefore call for broadly similar disclosure control protections as the release of the Jamaican national census dataset (both populations being approximately 3 million).

3. Results of comparative testing

In order to test the various post-tabular methods described above, they were applied to the cross-tabulations of census data from Jamaica, Belize and Saint Kitts and Nevis. The results for the census data from Belize appear in table 10 and are discussed below. The results for the census data from Jamaica and Saint Kitts and Nevis differed numerically, due to the different input data, but did not differ greatly in terms of what they revealed about the respective methods. Therefore, the results obtained using data from Jamaica and Saint Kitts and Nevis appear in tables A1 and A2 in Annex 1.

For the two-way cross-tabulation of age against sex, as would be expected, disclosure risks are low with this table having no small cells. All of the eight post-tabular methods successfully masked all the original cell counts.

The mean absolute distance between perturbed and original cell counts is a measure of data utility and is to a large extent predictable from the expected magnitude of the perturbations (given the original counts), differing only because the mean absolute deviation is calculated from a finite number of perturbations. A rough approximation for the mean absolute deviation is provided by the expected magnitude of the perturbations arising from the symmetric probability distribution which determines the perturbations for all except the small cell counts. Recalling figure 1, where the different methods of perturbation were ordered according to the variance of the perturbations, the mean absolute deviation will generally reflect this same ordering with the mean absolute deviation being roughly proportional to the standard deviation of the perturbations.

In the case of the three-way cross-tabulation (age by sex by ethnicity), for five out of the eight perturbations schemes, there are no disclosed cell counts. Not surprisingly, the highest proportion of disclosed cells was for the CP1 scheme for which $D = 1$ (1.76 per cent of cells disclosed or 16.5 per cent of small cells). There were smaller numbers of disclosed cells for the CP2 and RR3 schemes (for both of which $D = 2$), with no disclosed cells for the other five schemes, all $D \geq 3$.

For the four-way and five way cross-tabulations, five of the eight perturbation schemes continue to successfully mask all the original counts. These are increasingly sparse tabulations. In the case of the four-way tabulation, 77 per cent of cell values are either 0, 1 or 2, and in the case of the five-way tabulation 96 per cent. For these tabulations, the relative impact of the data perturbation is much greater, because the original cell values are much smaller, and this is reflected in the mean relative distance. Statistical offices will generally enforce some minimum threshold for the sparseness of the census tables that they publish, either formally or informally (see chapter IV section B). Of these two tables, the five-way table in particular would probably be deemed too sparse for publication (the mean cell size is 0.87).

Table 10
Measures of disclosure risk and information loss for post-tabular disclosure control methods applied to two-, three-, four- and five-way cross-tabulations of data from the Belize census, 2010

	Disclosed cells (percentages)	R(F,w ₁ ,w ₂)	Mean absolute distance	Mean relative distance	Hellinger's distance (range 0-1)	Small cells disclosed (percentages)	Mean absolute distance (small cells)
Two-way							
CP1	0	0.006	0.51	0	0		
CP2	0	0.005	0.7	0	0		
CP3(EU)	0	0.005	0.69	0	0		
RR3	0	0.004	0.88	0	0		
CP3	0	0.004	0.83	0	0		
CP4	0	0.004	0.98	0	0		
CP3(1,2)	0	0.004	1.01	0	0		
RR5	0	0.003	1.61	0	0		
Three-way							
CP1	1.76	0.053	0.5	0.07	0.01	16.49	0.52
CP2	0.05	0.041	0.74	0.11	0.01	0.43	0.78
CP3(EU)	0	0.041	0.73	0.11	0.01	0	0.78
RR3	0.21	0.033	0.92	0.15	0.01	1.95	1.31
CP3	0	0.034	0.87	0.12	0.01	0	0.89
CP4	0	0.029	1.03	0.14	0.01	0	0.96
CP3(1,2)	0	0.027	1.06	0.16	0.01	0	1.31
RR5	0	0.016	1.64	0.24	0.01	0	1.93
Four-way							
CP1	6.53	0.133	0.5	0.18	0.02	18.46	0.5
CP2	0.04	0.121	0.73	0.27	0.03	0.12	0.73
CP3(EU)	0	0.121	0.72	0.26	0.03	0	0.72
RR3	0.54	0.108	1.01	0.44	0.04	1.52	1.33
CP3	0	0.114	0.84	0.29	0.03	0	0.8
CP4	0	0.108	0.98	0.32	0.03	0	0.86
CP3(1,2)	0	0.103	1.12	0.46	0.04	0	1.33
RR5	0	0.083	1.71	0.64	0.05	0	1.86
Five-way							
CP1	8.5	0.221	0.5	0.23	0.03	18.43	0.5
CP2	0.08	0.214	0.73	0.34	0.04	0.17	0.73
CP3(EU)	0	0.214	0.73	0.33	0.04	0	0.72
RR3	0.83	0.204	1.05	0.57	0.06	1.81	1.34
CP3	0.01	0.21	0.84	0.37	0.04	0.01	0.8
CP4	0	0.206	0.96	0.39	0.05	0	0.86
CP3(1,2)	0	0.201	1.16	0.59	0.06	0.01	1.34
RR5	0	0.179	1.73	0.8	0.08	0.01	1.82

Source: ECLAC on the basis of Census 2010 data provided by the Statistical Institute of Belize.

Note: Statistics are calculated from all cells (interior cells and totals) except for R(F, w₁, w₂) which is calculated from interior cells only. Details of the test data appear in table 9. See figure 1 for details of the perturbation schemes.

Due to the ad-hoc nature of the disclosure risk measures for tabular data, it is difficult to argue that one of these methods is conclusively better than the others. The results show that the method CP1, applying maximum perturbations of ±1, does not provide adequate protection to serve as the primary method of disclosure control because the masking that it provides can frequently be 'unpicked'. Any of the other seven methods could plausibly be applied to provide disclosure control protection of census data, all of them providing slightly different levels of protection.

In practice, no method of disclosure control acts completely in isolation. Whatever method of disclosure control is chosen, the data producer still has to make decisions about what tables to release, how variables are coded, whether minimum population or other thresholds should apply. Therefore, the choice of perturbation method also depends on the context in which the method is being applied and the extent to which these other decisions about data release are also providing a certain measure of protection. For example, if a data producer supplements their use of post-tabular perturbation by some of these other protections, then

methods applying a more moderate degree of perturbation (CP₂, RR₃, CP₃(EU)) could be appropriate. If these other protections are either not in place, or are not so strong, then a greater degree of post-tabular perturbation is likely to be appropriate (CP₃, CP₄, CP_{3(1,2)}, RR₅).

While this analysis does not immediately demonstrate which of these methods is the best (and certainly not that one would be best in all circumstances), it does reveal something about the behaviour of the different methods and therefore can inform decisions about which methods are likely to be most useful to REDATAM users and what appropriate parameterizations of those methods might be in order to provide different levels of protection.

It should also be recognised that this analysis only considers disclosure risks associated with individual tables in isolation (albeit up to five dimensions). The possibility of database reconstruction attacks discussed in the previous chapter makes clear that disclosure risks arise not only from the possibility that an intruder is able to disclose information from a single cross-tabulation, but that they are able to combine information from multiple tables, perhaps a large number of tables, to partially reconstruct the microdata. Simulating a database reconstruction attack is an undertaking beyond the scope of what has been possible in this study. However, the existence of unquantified disclosure risks associated with more sophisticated forms of attack than it has been possible to recreate here, implies that the quantitative estimates of risk presented here underestimate the extent of what is possible, were a sufficiently motivated, skilled and well-resourced intruder to devote themselves to the task of a database reconstruction attack. All this suggests that in selecting from among the disclosure control methods analysed here, a degree of caution and allowance for some margin of safety is appropriate.

4. An additional feature of the cell key method

Analysing the way in which the additive relationships between cells can sometimes lead post-tabular perturbation methods to be ‘unpicked’ reveals an additional feature of the cell key method. One of the circumstances in which methods of cell perturbation and random rounding can be ‘unpicked’ for a marginal total and its corresponding interior cells, is when all the interior cells are perturbed in one direction to their maximum extent, and the corresponding marginal total is perturbed in the other direction to its maximum extent. Consider a simple example in which random perturbations of -1, 0 and +1 are applied to a multi-dimensional table, one dimension of which is sex. In this table there will be many sets of three cells which have the additive relationship: males plus females equals all persons. With any set of three cells, the two combinations of perturbations which can be unpicked are: males (-1), females (-1) and all persons (+1); and males (+1), females (+1) and all persons (-1). In each case, an intruder knowing nothing more than that ±1 are the maximum perturbations, can ‘unpick’ those perturbations to derive the original cell counts. This is because these two combinations of perturbations create additive discrepancies of 3 and -3 respectively between the two interior cells and the marginal total, and each of these combinations of perturbations is the only one that can create that discrepancy. An intruder can therefore work out what the perturbations are for these cells and recover the original counts. The same issue applies in theory to other cell perturbation schemes although with random perturbations of -3, -2, -1, 0, 1, 2 and 3, it is highly unlikely that this scheme would give rise to the perturbations: males (-3), females (-3) and all persons (+3), particularly since the more extreme perturbations are generally applied with the smallest probability. The issue also applies to random rounding, particularly to base three, and in this case, there can be no doubt that the intruder has knowledge of the method since it is obvious from the tables. However, it turns out that an additional feature of the cell key method is that it reduces the risk that original cell counts are disclosed in this way.

The traditional form of random rounding is to round marginal totals and the grand total independently of their corresponding interior cells. With the cell key method, rounding or perturbation of the totals is based on the corresponding cell key which is calculated as the sum (modulo 1) of the record keys for all records which are counted in that total cell. This means that the cell keys for the totals must also be equal to the sum (modulo 1) of the cell keys for the interior cells corresponding to each of those marginal totals. In other words, there is now some dependence between the perturbations applied to totals and the perturbations applied to the interior cells which correspond to each total.

This can be seen from a very simple example of a ‘ptable’ for the simplest perturbation scheme CP1 (see table 11). The ptable contains essentially the same information as the transition matrix but also illustrates how the cell keys are used to determine the perturbations. As explained above, the cell keys are calculated as the sum (modulo 1) of record keys uniformly distributed between 0 and 1. The cell keys also, therefore, take values between 0 and 1 and are compared with lower and upper bounds in the ptable to determine the perturbation. For an original count of 0, the cell key will always fall between the lower and upper bounds of 0 and 1 and there is only one possible perturbation which is zero. For an original count of 1 (and any other count), if the cell key falls in the range 0 to 0.25, the perturbation is -1; if the cell key falls in the range 0.25 to 0.75, the perturbation is 0; while if the cell key falls in the range 0.75 to 1, the perturbation is +1.

Table 11
Lookup table (or ‘ptable’) for the perturbation scheme CP1

Original count i	Perturbed count j	Probability p	Perturbation v	Lower bound lb	Upper bound ub
0	0	1	0	0	1
1	0	0.25	-1	0	0.25
1	1	0.5	0	0.25	0.75
1	2	0.25	1	0.75	1

Note: Produced using the R package ptable (Enderle, 2019).

Taking two interior cells summing to a marginal total (a male cell, a female cell and an ‘all persons’ cell), if the cell keys for both the male and female cells are in the range 0 to 0.25, leading to the perturbation -1 for both of them, then the only possible perturbations for the ‘all persons’ cell are either -1 or 0. This is because the cell key for ‘all persons’ must be equal to the sum (modulo 1) of the cell keys for males and females. If these both fall into the range 0 to 0.25, then the possible range of values for the cell key for ‘all persons’ is 0 to 0.5. Therefore, the cell key method prevents precisely the combination of perturbations—males (-1), females (-1) and all persons (+1)—which would disclose the original cell values. In the same way, if the male and female cells are both perturbed by +1, the cell key for all persons ensures that the only possible perturbations are 0 or +1, and the disclosive combination of males (+1), females (+1) and all persons (-1) is prevented.

This is a simple example which does not take account of the full multi-dimensional nature of the relationships between interior cells and marginal totals, however, it does help to explain the following results. Table 12 shows the percentage of disclosed cells for different perturbation schemes when totals and marginal totals are truly perturbed independently of the interior cells in the tables versus perturbation with the cell key. The cell key method consistently reduces the rate at which the additive relationships between cells can be used to ‘unpick’ the random perturbations or rounding adjustments and disclose original cell values.

Table 12
Perturbation of totals and marginal totals using the cell key: impact on disclosure risk
for cross-tabulations of data from the Saint Kitts and Nevis census, 2011
(Percentages)

	Disclosed cells		Small cells disclosed	
	Independent perturbation of totals	Perturbation of totals using the cell key	Independent perturbation of totals	Perturbation of totals using the cell key
Two-way				
CP1	2.04	0		
CP2	0	0		
CP3(EU)	0	0		
RR3	0	0		
CP3	0	0		
CP4	0	0		
CP3(1,2)	0	0		
RR5	0	0		

	Disclosed cells		Small cells disclosed	
	Independent perturbation of totals	Perturbation of totals using the cell key	Independent perturbation of totals	Perturbation of totals using the cell key
Three-way				
CP1	7.16	4.4	20.87	17.5
CP2	0.05	0	0.2	0
CP3(EU)	0	0	0	0
RR3	0.7	0.15	0.99	0.6
CP3	0	0	0	0
CP4	0	0	0	0
CP3(1,2)	0	0	0	0
RR5	0	0	0	0
Four-way				
CP1	11.37	6.33	26.93	19.14
CP2	0.25	0.07	0.43	0.21
CP3(EU)	0	0	0	0
RR3	1.18	0.69	2.56	2.09
CP3	0.02	0.01	0.06	0.02
CP4	0	0	0	0
CP3(1,2)	0.05	0	0.06	0
RR5	0	0	0	0
Five-way				
CP1	13.83	8.3	26.61	18.23
CP2	0.2	0.06	0.25	0.13
CP3(EU)	0	0	0	0
RR3	1.24	0.66	2.21	1.45
CP3	0.02	0	0.01	0
CP4	0	0	0	0
CP3(1,2)	0.05	0	0.02	0
RR5	0	0	0	0

Source: ECLAC on the basis of Census 2011 data provided by the Ministry of Sustainable Development of Saint Kitts and Nevis.

Note: Details of the test data appear in table 9. See figure 1 for details of the perturbation schemes.

IV. Implementation of disclosure control functionality in redatam

This report has thus far focused on the principal disclosure control methods employed by statistical offices to protect census tables. The statistical offices using these methods will generally apply them in conjunction with other disclosure control rules and thresholds which also play an important role in determining exactly what census tables are published. Some examples of such rules and thresholds include limiting the number of dimensions in output tables, minimum population thresholds and a minimum average cell size. In seeking to implement disclosure control functionality in REDATAM, consideration therefore needs to be given to the way in which REDATAM can provide not just a single method of disclosure control but also facilitate the use of accompanying rules and thresholds, either through existing functionality of REDATAM or by implementing new functionality.

A. Post-tabular perturbation (with cell keys)

Among Caribbean statistical offices, the most commonly used tools for final processing and tabulation of census results in the 2010 census round were CSPro, SPSS and REDATAM with most statistical offices using at least two or three of these software packages. The use of several software packages allows statistical offices to take advantage of the unique strengths of each, for example data editing in CSPro, data analysis in SPSS, and online dissemination through REDATAM (all this obviously has to be carefully managed). For the 2020 round of censuses, it is likely that some statistical offices will also be using Survey Solutions. REDATAM was used by 10 Caribbean statistical offices to disseminate results from the previous census round and its use is becoming more established among Caribbean statistical offices. Implementing post-tabular disclosure control in REDATAM would therefore be an effective way of making it available to statistical offices.

Post-tabular methods of disclosure control are most effective when they are applied to all census tables, including those in national census reports, in online dissemination through REDATAM, and in any other tabular outputs produced by the statistical office. If only some tables are protected but others are not, the protection of even those tables for which protection has been applied is weakened. This is because the release of true cell counts for some cells could be used to unpick the perturbation applied to other cells.

Cell key perturbation and random rounding could be made available in the Process Module of REDATAM to enable statistical offices to produce protected frequency tables. The method would also need to be available in the Webserver module of REDATAM so that the online applications for future censuses could also incorporate the method. A range of perturbation and random rounding schemes could be offered so that users could choose between relatively lighter or stronger perturbation of the data according to their national situation.

B. Complementary methods

Sparse tables with many zeros and small counts present the highest disclosure risk and they are also the tables upon which cell perturbation will have the greatest relative impact. For a sparse table, the average absolute distance between the cell values of the original and perturbed table will be similar to that for a less sparse table. The average relative distance, on the other hand, will be much larger for a sparse table because the perturbations themselves are of a similar order of magnitude to many of the cell counts. There are, thus, good reasons for restricting the release of very sparse tables, both to protect confidentiality and for reasons of data quality.

Several national statistical offices have applied a minimum threshold for the average cell size and suppressed tables which fall below this threshold. In their 2018 census, Statistics New Zealand used a minimum mean cell size of 2 which is applied separately for each geographical unit (SNZ, 2019). The United Kingdom applied a mean cell size threshold of 1 in their 2001 and 2011 censuses (ONS, 2004 and 2012). In their 2000 and 2010 censuses, the US Census Bureau also made some use of a mean cell size threshold of 3 for special tabulations (McKenna, 2018).

The possibility to apply a mean cell size threshold to suppress sparse tables could be useful in REDATAM, particularly for online applications. This could be applied in a couple of ways. Firstly, the online application could simply suppress tables when they fall below a certain minimum mean cell size. An alternative, implemented by Statistics New Zealand, is to use the mean cell size threshold to identify 'sensitive' sparse tables and then apply cell suppression to small cell counts. This cell suppression can be applied in a straightforward way, simply replacing all small cell values with (c) to signify confidential. In this case small cells are those less than 6, which in practice means cells taking the value 0 or 3, remembering that Statistics New Zealand apply random rounding to base 3. No secondary cell suppression is necessary because of the random rounding and non-additivity of the data. This approach has the advantage that tables are not suppressed in their entirety and users are provided with at least some information.

C. Disclosure control protection for continuous variables

This report has focused on protecting frequency tables of census data because the vast majority of tabulations of Caribbean census data consist of frequency counts. There are a much smaller number of continuous variables collected in censuses, for example income in some countries. Even income is often collected in bands or otherwise converted from a continuous to an ordinal categorical variable by grouping the data (with concerns about data accuracy sometimes cited as a reason for banding the data). In this case, income will be tabulated in frequency tables in the same way as for other categorical variables. However, for any genuine continuous variables collected in censuses, they are tabulated in what are referred to as magnitude tables. Magnitude tables are used to represent the relationship between one continuous variable and one or more categorical variables, with table cells defined by the categorical variables and cell values calculated as the average of the values on the continuous variable for the cases falling into each cell. An example of a magnitude table might be average earnings by occupation and sex. Disclosure risks arise when the cell value is the average of a small number of cases, especially 1 or 2 cases. In this case, an intruder's knowledge of an individual may enable them to identify contributors and therefore discover information about their earnings.

The ABS also developed a variant of the cell perturbation method that can also be used to protect magnitude tables showing averages of continuous variables (Thompson, Broadfoot and Elazar, 2013). A variant of this method has also been implemented in the CellKey package (Giessing and Tent, 2019; and Meindl and Enderle, 2019). Thompson, Broadfoot and Elazar describe this as the Top Contributors Method which consists of pseudo-random multiplicative adjustments made to the top contributors of each non-zero cell of the table. Tambay (2017) proposed the Layered Perturbation Method which is based on similar ideas and incorporates a noise specification designed specifically to target disclosure by differencing. REDATAM produces magnitude tables using the 'average' process. If the cell perturbation method and random rounding were implemented in the REDATAM software, a version of the Top Contributors or Layered Perturbation Method could also be implemented so that methods for protecting both frequency and magnitude tables were available.

D. REDATAM's programming window

REDATAM web applications can provide users with the functionality to specify queries using REDATAM's programming language. The language consists of basic commands for tabulating data and computing indicators. However, REDATAM applications implemented in the last few years for Caribbean census data have not offered this functionality precisely because of concerns about disclosure risk.

This study has proposed the use of methods which could also be integrated into the REDATAM programming window which would strengthen protection of confidentiality for relatively little impact on data utility. All outputs returned to users of the programming language are in the form of tables. This means that the post-tabular perturbation methods such as those discussed in this study could be applied, by default, to tabular outputs returned to users through the programming window. Other complementary measures such as those discussed above could also be applied.

However, privacy protection in a context where users have programming functionality to specify queries is a rather different problem compared with confidentiality protection for a menu driven table generator. One of the main differences is that with menu driven functionality, applications can be designed in such a way that some limits are placed on the range and volume of tabulations that are available to users (for example limiting which variables can be cross-tabulated). Recent work on database reconstruction suggests that such constraints are important because if a sufficiently large range and volume of tabulations from a confidential data source are made public, then it becomes possible to reconstruct the confidential database (see chapter II, section C). The risk here is from sophisticated attacks because database reconstruction is not a trivial matter. The addition of noise certainly makes these attacks more difficult, but this raises another important trade-off between the volume of statistics that are made available and the amount of noise that is necessary to protect those tabulations. The more tabulations that are made available, the more noise is required to protect those tables, and at some point, it becomes preferable to limit the volume of tabulations in order to be able to limit the amount of noise that is applied to the data. With a programming window, it is much less clear how to restrict the volume of tabulations that are available. Other parallel measures which address the problem from a slightly different angle could include monitoring mechanisms designed to first assess that users are human beings and, secondly, whether their use of REDATAM applications was consistent with plausible and legitimate research use. Registration and vetting of users could also be considered. However, these are areas where further work is required before concrete steps forward can be considered.

V. Conclusions

The use of post-tabular methods of disclosure control would enable Caribbean statistical offices to publish and disseminate richer, more detailed and disaggregated statistics than can safely be released at present. The most suitable methods of disclosure control for Caribbean censuses are post-tabular methods, either cell perturbation or random rounding, which introduce a small amount of additive noise to the original frequency counts. Giving up a very small amount of accuracy would be a price worth paying to enable the release of more disaggregated data than is currently published.

These methods are relatively simple to implement and can provide strong protection of statistical confidentiality. The REDATAM software, which is primarily a tabulation and dissemination tool, provides the perfect vehicle for implementing post-tabular disclosure control methods as an add-on to REDATAM's tabulation functionality. If appropriate routines were built into REDATAM, their adoption ought to be relatively easy.

Pre-tabular methods of disclosure control are not so well-suited to Caribbean censuses. They do not provide such strong protection of confidentiality, particularly for flexible table generators; the most widely used pre-tabular method, record swapping, is not applicable to most Caribbean countries due to their small size; and implementation of pre-tabular methods would introduce a new and relatively complicated additional step into census data processing. While routines implementing post-tabular methods could be made available to statistical offices through the REDATAM software, the same cannot be said of pre-tabular methods which would be a much less natural fit with REDATAM.

Neither is the differential privacy approach suitable for the Caribbean. At the time of writing, no national statistical office has released differentially private census results. The United States Census Bureau will do so for the first time when the results of the 2020 census are published, but those methods which are being used to protect the confidentiality of the United States' short form census are not applicable to the long form censuses used in the Latin America and Caribbean region.

This leaves the cell perturbation and random rounding methods both of which provide effective disclosure control protection for census tables. They are flexible to the extent that they can be set to apply slightly more or

less strong protection of the data, according to the user's needs. There is relatively little to choose between the cell perturbation and random rounding methods and both could be implemented in REDATAM.

As discussed in this study, one clear disadvantage of both cell perturbation and random rounding is that tables are no longer strictly additive. This is certainly a feature of the methods that some users, and some statisticians find unappealing. Producers and users of statistics are generally accustomed to statistical tables in which there are strict additive relationships within tables.

Statistical offices adopting new approaches to statistical disclosure control for the 2020 round of censuses also take different positions on this issue. As noted in this study, in communicating their decision to shift from a pre-tabular to a post-tabular method of disclosure control for their 2021 census, the German Federal Statistical Office described non-additivity as "not much more than a cosmetic flaw" (Destatis, 2018). The United Kingdom's ONS, by contrast, reported that prior to the 2011 census, additivity and consistency had been users "uppermost concern." More recent feedback did suggest that for future censuses, users would value "flexibility, accessibility and timeliness, above additivity and consistency," but the ONS retained a preference for presenting each individual table as internally additive (thus the hybrid approach combining pre- and post-tabular methods of disclosure control) (Spicer, 2020).

Statistical offices that have published census data in this way, in some cases for multiple censuses, report that while data users regard additivity of tables as a desirable property, they prefer non-additivity to greater suppression of data (Stats NZ, 2019). It should also be remembered that non-additivity of statistical outputs is far from unheard of. Chain-linked volume measures of national accounts aggregates are non-additive. Other statistical processes such as conventional deterministic rounding and seasonal adjustment can also introduce non-additivity depending on how they are applied.

The application of post-tabular disclosure control could help statistical offices in both the Caribbean and Latin America as well as REDATAM users outside the region. If there are countries that are under-protecting their census data, these methods would reduce the risk of any breach of confidentiality. In countries where concerns about confidentiality have led to very restrictive release practices, the use of such methods could encourage government statisticians to make more disaggregated data available. In all these countries, judicious use of perturbative methods could enable census statisticians to strike a better balance between disclosure risk and data utility in dissemination of their census results.

Bibliography

- ABS (Australian Bureau of Statistics) (2019) "ABS response to articles about potential vulnerability in ABS TableBuilder tool", [online], Canberra, Australia [date of reference: 18 October 2020] <<https://www.abs.gov.au/websitedbs/d3310114.nsf/home/ABS%20response%20to%20articles%20about%20potential%20vulnerability%20in%20ABS%20TableBuilder%20tool>>.
- Antal, Laszlo, Tobias Enderle and Sarah Giessing (2017), "Statistical disclosure control methods for harmonised protection of census data", Work Package 3, Deliverable D3.1 Part I, Harmonised protection of census data in the ESS [online], [date of reference: 30 May 2020] <https://ec.europa.eu/eurostat/cros/content/methods-protecting-cesus-data_en>.
- Asghar, Hassan Jameel and Dali Kaafar (2020), "Averaging Attacks on Bounded Noise-based Disclosure Control Algorithms", *Proceedings on Privacy Enhancing Technologies* Vol. 2020 Issue 2, Sciendo, April.
- Bailie, James and Joseph Chien (2019) "ABS Perturbation Methodology Through the Lens of Differential Privacy", presented at the Joint UNECE/Eurostat work session on statistical data confidentiality, The Hague, the Netherlands, October.
- Blanchard, Stephanie (2019), "The methodological challenges of protecting outputs from a flexible dissemination system", *Survey Methodology Bulletin* No. 79, London, United Kingdom, Office for National Statistics, (ONS), May.
- Buzzigoli, Lucia and Antonio Giusti (1999), "An algorithm to calculate the lower and upper bounds of the elements of an array given its marginals", *Statistical Data Protection (SDP 1998, Lisbon, Portugal, March) Proceedings of the Conference*, Luxembourg, Eurostat.
- Chipperfield, James, Daniel Gow and Bronwyn Loong (2016), "The Australian Bureau of Statistics and releasing frequency tables via a remote server", *Statistical Journal of the IAOS* Vol. 32 Issue 1, International Association for Official Statistics.
- Cleveland, Lara and others (2012), "When Excessive Perturbation Goes Wrong and Why IPUMS International Relies Instead on Sampling, Suppression, Swapping, and Other Minimally Harmful Methods to Protect Privacy of Census Microdata", *Privacy in Statistical Databases 2012*, Josep Domingo-Ferrer and Ilaria Tinnirello (eds.), Palermo, Sicily, September, Proceedings, Springer.

- Destatis (2018), "Confidentiality in the 2021 Census", *Methods – Approaches – Developments: Information of the German Federal Statistical Office*, Edition 2/2018, Wiesbaden, Germany, December.
- Dove, Iain, Stephanie Blanchard, and Keith Spicer (2017), "Applying Cell-Key Perturbation to 2021 Census Outputs" [online], [date of reference: 22 June 2020] <<https://gss.civilservice.gov.uk/wp-content/uploads/2017/01/ExN-Disclosure-control-methodology-in-2021-Census-outputs-Spicer-Blanchard-Dove-ONS.docx>>.
- Duncan, George, Mark Elliot and Juan-José Salazar-González (2011), "Statistical Confidentiality - Principles and Practice", *Statistics for Social and Behavioral Sciences*, Springer.
- Enderle, Tobias (2019), "ptable: A Perturbation Table Generator for SDC Tools" [online], [date of reference: 05 June 2020] <<https://github.com/tenderle/ptable>>.
- Enderle, Tobias, Sarah Giessing, and Reinhard Tent (2018), "Designing Confidentiality on the Fly Methodology – Three Aspects", *Privacy in Statistical Databases 2018*, Josep Domingo-Ferrer and Francisco Montes (eds.), Valencia, Spain, September, Proceedings, Springer.
- Eurostat (2020), "Harmonised protection of Census data" [online], [date of reference: 12 May 2020] <https://ec.europa.eu/eurostat/cros/content/harmonised-protection-census-data_en>.
- Fraser, Bruce and Janice Wooten (2005), "A Proposed Method for Confidentialising Tabular Output to Protect against Differencing", presented at the Joint UNECE/Eurostat work session on statistical data confidentiality, Geneva, Switzerland, November.
- Garfinkel, Simson (2019), "Formal Privacy: Making an Impact at Large Organizations", presented at 2019 Joint Statistical Meetings, American Statistical Association, Denver, Colorado, July.
- _____(2017), "Modernizing disclosure avoidance: report on the 2020 disclosure avoidance subsystem as implemented for the 2018 end-to-end test (Continued)", presented at the 2017 Census Scientific Advisory Committee Fall Meeting, Maryland, United States of America, September.
- Garfinkel, Simson, John Abowd and Sarah Powazek (2018), "Issues encountered deploying differential privacy", Proceedings of the 2018 Workshop on Privacy in the Electronic Society, New York, United States of America, Association for Computing Machinery, January.
- Giessing, Sarah (2016), "Computational Issues in the Design of Transition Probabilities and Disclosure Risk Estimation for Additive Noise", *Privacy in Statistical Databases 2016*, Josep Domingo-Ferrer and Mirjana Pejic-Bach (eds.), Dubrovnik, Croatia, September, Proceedings, Springer.
- Giessing, Sarah and Jörg Höhne (2011), "Eliminating small cells from census counts tables: empirical vs. design transition probabilities", *SORT (Statistics and Operations Research Transactions)*, Special issue: Privacy in statistical databases, Statistical Institute of Catalonia, Spain, September.
- _____(2010), "Eliminating Small Cells from Census Counts Tables: Some Considerations on Transition Probabilities", *Privacy in Statistical Databases 2010*, Josep Domingo-Ferrer and Emmanouil Magkos (eds.), Corfu, Greece, September, Proceedings, Springer.
- Giessing, Sarah and Reinhard Tent (2019), "Concepts for generalising tools implementing the cell key method to the case of continuous variables", presented at the Joint UNECE/Eurostat work session on statistical data confidentiality, The Hague, the Netherlands, October.
- Hundepool, Anco and others (2012), "Statistical Disclosure Control", Wiley Series in Survey Methodology, John Wiley & Sons Ltd., United Kingdom.
- _____(2010), "Handbook on Statistical Control Version 1.2", ESSNetSDC (A Network of Excellence in the European Statistical System in the field of Statistical Disclosure Control).
- Jones, Francis and Kristin Fox (2016), "Dissemination of Caribbean census microdata to researchers: Including an experiment in the anonymization of microdata for Grenada and Trinidad and Tobago", *Studies and Perspectives Series – The Caribbean*, No.49 (LC/L.4134), Santiago, Chile, Economic Commission for Latin America and the Caribbean (ECLAC), January.
- Leaver, Victoria (2009), "Implementing a method for automatically protecting user-defined census tables", presented at the Joint UNECE/Eurostat work session on statistical data confidentiality, Bilbao, Spain, December.

- Leclerc, Philip (2019), "Results from a consolidated database reconstruction and intruder re-identification attack on the 2010 decennial census", presented at the workshop "Challenges and New Approaches for Protecting Privacy in Federal Statistical Programs", June.
- March, Mary and Douglas Norris (1987), "Disclosure avoidance techniques in the Canadian censuses of population and agriculture", Proceedings of the Survey Research Methods Section American Statistical Association.
- McKenna, Laura (2018), "Disclosure Avoidance Techniques Used for the 1970 through 2010 Decennial Censuses of Population and Housing", Washington, United States Census Bureau, November.
- Meindl, Bernhard (2020), "Introduction to the cellKey-Package" [online], [date of reference: 03 May 2020] <<https://sdctools.github.io/cellKey/articles/introduction.html>>.
- Meindl, Bernhard and Tobias Enderle (2019), "cellKey - consistent perturbation of statistical tables", presented at the Joint UNECE/Eurostat work session on statistical data confidentiality, The Hague, the Netherlands, October.
- National Statistics (2004), "Census 2001 Definitions", London, United Kingdom,
- Page, Hector, Charlie Cabot, and Kobbi Nissim (2018), "Differential privacy: an introduction for statistical agencies", A contributing article to the National Statistician's Quality Review into Privacy and Data Confidentiality Methods, United Kingdom, Government Statistical Service, December.
- Rinott, Yosef and others (2018), "Confidentiality and differential privacy in the dissemination of frequency tables", *Statistical Science* Vol. 33 No. 3, Cleveland, United States of America, Institute of Mathematical Statistics, August.
- Shlomo, Natalie (2018), "Methods to assess and quantify disclosure risk and information loss under statistical disclosure control", a contributing article to the National Statistician's Quality Review into Privacy and Data Confidentiality Methods, United Kingdom, December.
- _____(2007), "Statistical Disclosure Control Methods for Census Frequency Tables", *International Statistical Review* Vol. 75, No. 2, International Statistical Institute, August 2007.
- Shlomo, Natalie and Caroline Young (2008), "Invariant Post-tabular Protection of Census Frequency Counts", *Privacy in Statistical Databases 2008*, Josep Domingo-Ferrer and Yücel Saygin (eds.), Istanbul, Turkey, September, Proceedings, Springer.
- _____(2005), "Information loss measures for frequency tables", presented at the Joint UNECE/Eurostat work session on statistical data confidentiality, Geneva, Switzerland, November.
- Shlomo, Natalie, Laszlo Antal and Mark Elliot (2015), "Measuring Disclosure Risk and Data Utility for Flexible Table Generators", *Journal of Official Statistics*, Vol. 31, No. 2, Stockholm, Sweden, Statistics Sweden, June.
- _____(2013), "Disclosure Risk and Data Utility in Flexible Table Generators", presented at NTTS (New Techniques and Technologies for Statistics) 2013, Brussels, Belgian, March.
- Spicer, Keith (2020), "Statistical Disclosure Control (SDC) for 2021 UK Census" [online], London, United Kingdom, UK Statistics Authority, [date of reference: 1 August 2020] <<https://uksa.statisticsauthority.gov.uk/wp-content/uploads/2020/07/EAP125-Statistical-Disclosure-Control-SDC-for-2021-UK-Census.docx>>.
- STATBEL (The Belgian Statistical Office) (2015), "La protection des données grâce au swapping" [online], Brussels, Belgian [date of reference: 20 September 2020] <https://www.census2011.be/doc/census_workshop_swapping_fr.pdf>.
- Statistik Austria (n/d), "Target Record Swapping" [online], Vienna, Austria [date of reference: 21 September 2020] <http://www.statistik.at/web_de/static/record_swapping_063399.pdf>.
- Stats NZ (Statistics New Zealand) (2019), "Applying confidentiality rules to 2018 Census data and summary of changes since 2013", Wellington, New Zealand, September.
- Tambay, Jean-Louis (2017), "A layered perturbation method for the protection of tabular outputs", *Survey Methodology*, Vol. 43, No. 1, Ottawa, Canada, Statistics Canada, June.

- Thompson, Gwenda, Stephen Broadfoot and Daniel Elazar (2013), "Methodology for the Automatic Confidentialisation of Statistical Outputs from Remote Servers at the Australian Bureau of Statistics", presented at the Joint UNECE/Eurostat work session on statistical data confidentiality, Ottawa, Canada, October.
- USCB (United States Census Bureau) (2011), "2010 Census Redistricting Data (Public Law 94-171) Summary File" [online], Washington, [date of reference: 31 July 2020] <<https://www.census.gov/prod/cen2010/doc/pl94-171.pdf>>.
- Wood and others (2018), "Differential Privacy: A Primer for a Non-Technical Audience", *Vanderbilt Journal of Entertainment & Technology Law*, Vol. 21, No. 17, Berkman Klein Center Research Publication No. 2019-2.
- Zeisset, Paul (1978), "Suppression vs. Random rounding: disclosure-avoidance alternatives for the 1980 census", United States Census Bureau Working Papers [online], [date of reference: 20 September 2020] <<https://www.census.gov/library/working-papers/1978/adrm/cdar1978-01.html>>.

Annexes

Annex 1

Further results of comparative testing

Table A1
Measures of disclosure risk and information loss for post-tabular disclosure control methods applied to two-, three-, four- and five-way cross-tabulations of data from the Jamaica census, 2011

	Disclosed cells (percentages)	R(F,w1,w2)	Mean absolute distance	Mean relative distance	Hellinger's distance (range 0-1)	Small cells disclosed (percentages)	Mean absolute distance (small cells)
Two-way							
CP1	0	0.004	0.46	0	0		
CP2	0	0.003	0.7	0	0		
CP3(EU)	0	0.003	0.68	0	0		
RR3	0	0.002	0.88	0	0		
CP3	0	0.003	0.79	0	0		
CP4	0	0.002	0.92	0	0		
CP3(1,2)	0	0.002	0.96	0	0		
RR5	0	0.001	1.69	0	0		
Three-way							
CP1	0.49	0.043	0.5	0.03	0	20.41	0.53
CP2	0	0.032	0.72	0.04	0	0	0.82
CP3(EU)	0	0.032	0.71	0.04	0	0	0.76
RR3	0.1	0.028	0.88	0.05	0	4.08	1.16
CP3	0	0.026	0.85	0.05	0	0	0.9
CP4	0	0.022	1.02	0.05	0	0	0.96
CP3(1,2)	0	0.022	1.05	0.06	0	0	1.16
RR5	0	0.013	1.61	0.09	0	0	1.8
Four-way							
CP1	4.03	0.094	0.5	0.12	0	18.08	0.51
CP2	0.03	0.079	0.74	0.17	0.01	0.14	0.74
CP3(EU)	0	0.079	0.73	0.17	0.01	0	0.73
RR3	0.32	0.066	0.96	0.27	0.01	1.45	1.33
CP3	0	0.071	0.86	0.18	0.01	0	0.81
CP4	0	0.065	1.01	0.2	0.01	0	0.87
CP3(1,2)	0	0.06	1.1	0.29	0.01	0	1.33
RR5	0	0.042	1.68	0.4	0.01	0	1.85
Five-way							
CP1	5.25	0.189	0.5	0.16	0.01	17.82	0.5
CP2	0.05	0.178	0.73	0.23	0.01	0.17	0.73
CP3(EU)	0	0.178	0.73	0.23	0.01	0	0.72
RR3	0.56	0.167	0.98	0.37	0.02	1.87	1.33
CP3	0	0.171	0.85	0.25	0.01	0.01	0.8
CP4	0	0.164	0.99	0.28	0.02	0	0.87
CP3(1,2)	0.01	0.161	1.11	0.39	0.02	0.03	1.33
RR5	0	0.136	1.69	0.55	0.03	0.01	1.87

Source: ECLAC on the basis of Census 2011 data provided by the Statistical Institute of Jamaica.

Note: Statistics are calculated from all cells (interior cells and totals) except for R(F, w₁, w₂) which is calculated from interior cells only.

Details of the test data appear in table 9. See figure 1 for details of the perturbation schemes.

Table A2
Measures of disclosure risk and information loss for post-tabular disclosure control methods applied to two-, three-, four- and five-way cross-tabulations of data from the Saint Kitts and Nevis census, 2011

	Disclosed cells (percentages)	R(F,w ₁ ,w ₂)	Mean absolute distance	Mean relative distance	Hellinger's distance (range 0-1)	Small cells disclosed (percentages)	Mean absolute distance (small cells)
Two-way							
CP1	0	0.009	0.5	0.01	0		
CP2	0	0.008	0.69	0.01	0		
CP3(EU)	0	0.008	0.68	0.01	0		
RR3	0	0.007	0.93	0.01	0		
CP3	0	0.007	0.81	0.01	0		
CP4	0	0.006	0.98	0.01	0		
CP3(1,2)	0	0.006	1.01	0.01	0		
RR5	0	0.005	1.59	0.02	0		
Three-way							
		0.169					
CP1	4.4	0.141	0.5	0.13	0.01	17.5	0.47
CP2	0	0.125	0.73	0.19	0.02	0	0.68
CP3(EU)	0	0.125	0.71	0.18	0.02	0	0.65
RR3	0.15	0.109	0.96	0.32	0.02	0.6	1.29
CP3	0	0.114	0.84	0.21	0.02	0	0.74
CP4	0	0.106	0.98	0.22	0.02	0	0.8
CP3(1,2)	0	0.101	1.1	0.34	0.03	0	1.29
RR5	0	0.074	1.7	0.5	0.04	0	1.92
Four-way							
		0.242					
CP1	6.33	0.228	0.49	0.17	0.03	19.14	0.49
CP2	0.07	0.217	0.72	0.25	0.04	0.21	0.72
CP3(EU)	0	0.217	0.7	0.24	0.04	0	0.69
RR3	0.69	0.204	0.99	0.42	0.05	2.09	1.32
CP3	0.01	0.21	0.82	0.27	0.04	0.02	0.77
CP4	0	0.203	0.95	0.29	0.04	0	0.83
CP3(1,2)	0	0.198	1.11	0.44	0.06	0	1.32
RR5	0	0.169	1.7	0.62	0.07	0	1.85
Five-way							
		0.301					
CP1	8.3	0.296	0.49	0.23	0.04	18.23	0.49
CP2	0.06	0.292	0.73	0.34	0.05	0.13	0.73
CP3(EU)	0	0.292	0.72	0.33	0.05	0	0.72
RR3	0.66	0.285	1.04	0.57	0.08	1.45	1.33
CP3	0	0.289	0.83	0.36	0.06	0	0.79
CP4	0	0.286	0.95	0.39	0.06	0	0.84
CP3(1,2)	0	0.283	1.16	0.59	0.08	0	1.33
RR5	0	0.265	1.7	0.79	0.1	0	1.78

Source: ECLAC on the basis of Census 2011 data provided by the Ministry of Sustainable Development of Saint Kitts and Nevis.

Note: Statistics are calculated from all cells (interior cells and totals) except for R(F, w₁, w₂) which is calculated from interior cells only.

Details of the test data appear in table 9. See figure 1 for details of the perturbation schemes.

Annex 2

Technical annex

This annex provides some additional technical details of the methods and diagnostics used in this study.

The cell key perturbation method

The version of the cell key perturbation method used here was that implemented in the R package `cellKey` (Meindl and Enderle, 2019). It differs slightly from the original ABS version. In the `cellKey` version, record keys are randomly sampled from a uniform distribution over the range zero to one and assigned to every record in the database (e.g. person, household, building, dwelling).

$$rkey \in (0,1]$$

Cell keys are then derived as the sum (modulo 1) of the record keys for the n records in each cell:

$$ckey = \sum_{i=1}^n rkey_i \bmod 1$$

All cell keys therefore take values in the same range as the record keys: (0,1]. Table A3 illustrates how the 'ptable' uses the cell key to perturb the original cell counts.

Table A3
Lookup table (or 'ptable') equivalent to transition matrix B in diagram 1 (first 13 rows only)

i	j	p	v	lb	ub
0	0	1.00000000	0	0.00000000	1.00000000
1	0	0.66750000	-1	0.00000000	0.66750000
1	3	0.33000000	2	0.66750000	0.99750000
1	4	0.00250000	3	0.99750000	1.00000000
2	0	0.33416658	-2	0.00000000	0.33416658
2	3	0.66333419	1	0.33416658	0.99750077
2	4	0.00249871	2	0.99750077	0.99999948
2	5	0.00000052	3	0.99999948	1.00000000
3	0	0.14853924	-3	0.00000000	0.14853924
3	3	0.50678866	0	0.14853924	0.65532790
3	4	0.25654543	1	0.65532790	0.91187333
3	5	0.07530772	2	0.91187333	0.98718105
3	6	0.01281895	3	0.98718105	1.00000000

Note: Produced using the R package ptable (Enderle, 2019).

The columns of the ptable are as follows:

- i: defines a "block" in which perturbations for original counts of i can be found
- j: possible perturbed values for given values of i
- p: the probability that a particular perturbation is selected for a given original cell count i
- v: the perturbation value

The columns lb and ub are calculated directly from p and the perturbation for an original cell count of i is determined by comparing the cell key with lb and ub. From the first row of the table, an original cell count of $i = 0$ becomes $j = 0$ with probability 1 because the cell key must fall in the range 0 to 1. From row 2, an original cell count of $i = 1$ becomes $j = 0$ if the cell key falls within the range 0 to 0.6675 which occurs with probability 0.6675.

Measures of risk

The proportion of disclosed cells is calculated using the approach described in Enderle, Giessing and Tent (2018) and Giessing (2016), albeit with a couple of variations. Whereas in the original implementation, the disclosure scenario assumed all original non-zero cells to be unknown by the intruder, in this case all original cells were assumed to be unknown including zeros, reflecting the fact that some of the zeros in the perturbed table were non-zero in the original table. Disclosed cells are deemed to be those where the algorithm computes upper and lower bounds for cell values which are equal or cells where the lower bound is zero and the upper bound is one. This latter case is deemed to constitute disclosure since an intruder is assumed to be trying to link the cell to an individual that exists and therefore the intruder possesses the information to rule out the possibility that the true cell value is zero. A further important difference concerned the test data which in the original implementation were EU census hypercubes which have a more complex hierarchical structure than the cross-tabulations produced using REDATAM. For example, in the EU census hypercube used by Enderle, Giessing and Tent, the age variable is single year of age but includes five-yearly subtotals, whereas REDATAM tables simply contain a single total for each variable. In an analysis of the extent to which additive table relationships can be used to ‘unpick’ post-tabular perturbation methods, clearly the more additive relationships that a table has, the more potential there is for an intruder to use them to derive original cell values.

The other measure of risk is that adapted from Shlomo, Antal and Elliot (2015) where the risk measure for unperturbed tables is:

$$R(F, w_1, w_2) = w_1 \cdot \left[\frac{|A|}{K} \right] + w_2 \cdot \left[1 - \frac{N \cdot \log N - \sum_{i=1}^K F_i \cdot \log F_i}{N \cdot \log K} \right] \\ - (1 - w_1 - w_2) \cdot \left[\frac{1}{\sqrt{N}} \cdot \log \frac{1}{e\sqrt{N}} \right]$$

where $|A|$ is the number of zeroes in the interior cells of the table, K is the number of interior cells, F is the vector of the frequency counts in those cells $F = (F_1, F_2, \dots, F_K)$, and $N = \sum_{i=1}^K F_i$ while w_1 and w_2 are arbitrary weights: $0 \leq w_1 + w_2 \leq 1$.

The first term of this composite indicator of disclosure risk is a measure of the proportion of zeros in the interior cells of a table. A high number of zeros creates disclosure risk because it increases the likelihood of having rows or columns with only one non-zero cell (see chapter I on group disclosure).

The second component indicator of risk is a measure of entropy, in other words, the extent to which interior cell counts are evenly spread across the different categories of the variables forming the table, versus the extent to which they are clustered in one or a small number of cross-classified categories. High entropy, or spread, of the frequency counts across a table means greater uncertainty and difficulty for an intruder in trying to identify a known individual in a particular cell of a table. This entropy term is actually split and is calculated separately for small cell counts (0 to 6) and large cell counts (>6) so that the small cell counts can be allocated a greater weight, reflecting their importance in the context of disclosure risk. The third indicator of risk is a function of N , the total population for that table. The larger the population, the lower the disclosure risk.

These three indicators of risk are weighted together using the same weighting scheme as originally proposed based on analysis of census data from the United Kingdom: $w_1 = 0.1$ (zero cells), $w_2Part1 = 0.7$ and $w_2Part2 = 0.1$ (entropy), $w_3 = 0.1$ (total population). This measure of disclosure risk can be calculated for the original table. When it is calculated for the perturbed table, the first and second term are modified to account for the uncertainty which is introduced by the post-tabular disclosure control. The first term is adjusted to account for additional false zeros introduced by the data perturbation (original cell counts greater than zero which become zero after perturbation). In a

simplification of the originally proposed measure, the entropy terms are discounted by a factor $1/(1 + \text{var}(v))$, where $\text{var}(v)$ is the variance of the random perturbations introduced to the cell counts. So, the greater the variance of the perturbations, the lower the disclosure risk. This measure of risk is denoted⁵ $R(F, w_1, w_2)$ since it is a function of the frequency counts ($F = (F_1, F_2, \dots, F_K)$) and the weights (with $w_3 = 1 - w_1 - w_2$). It is bounded by 0 (low risk) and 1 (high risk).

Measures of data utility

Using the same notation as Chipperfield, Gow and Loong (2016), tables consist of T cells indexed $t = 1, \dots, T$ which include both interior cells, a total and marginal totals. The cell values in the original table are denoted c_t while the cell values of the perturbed table are denoted m_t (denoting the fact that the true cell values have been masked).

The following measures of data utility were calculated for each different perturbation scheme.

The mean absolute distance (per cell):
$$\frac{1}{T} \sum_{t=1}^T |m_t - c_t|$$

The mean relative absolute difference (per cell):
$$\frac{1}{T} \sum_{t=1}^T \frac{|m_t - c_t|}{c_t}$$

Hellinger distance:
$$\sqrt{\frac{1}{2} \sum_{t=1}^T (\sqrt{m_t} - \sqrt{c_t})^2}$$

⁵ The original notation used by Shlomo, Antal and Elliot (2015) ($R(F, w_1, w_2)$) has been retained here although, as mentioned above, the calculation of the entropy terms for the post perturbation tables was somewhat different.



UNITED NATIONS

Series

E C L A C

Studies and Perspectives-The Caribbean

Issues published

A complete list as well as pdf files are available at
www.eclac.org/publicaciones

- 96. Statistical disclosure control for Caribbean census tables: a proposal to expand the availability of disaggregated census data, Francis Jones, (LC/TS.2020/159, LC/CAR/TS.2020/6), 2020.
- 95. Economic implications of the ban on single-use plastics in the Caribbean: a case study of Trinidad and Tobago, Willard Phillips, Elizabeth Thorne and Camille Roopnarine (LC/TS.2020/127, LC/CAR/TS.2020/5), 2020.
- 94. Creating an enabling environment for e-government and the protection of privacy rights in the Caribbean: a review of data protection legislation for alignment with the General Data Protection Regulation, Amelia Bleeker (LC/TS.2020/126, LC/CAR/TS.2020/4), 2020.
- 93. The use of technology and innovative approaches in disaster and risk management: a characterization of Caribbean countries' experience, Luciana Fontes de Meira, Omar Bello (LC/TS.2020/106, LC/CAR/TS.2020/3), 2020.
- 92. Preliminary overview of the economies of the Caribbean 2019–2020, Dillon Alleyne, Michael Hendrickson, Sheldon McLean, Maharouf Oyolola, Machel Pantin, Nyasha Skerrette and Hidenobu Tokuda (LC/TS.2020/56, LC/CAR/TS.2020/2), 2020.
- 91. Caribbean synthesis report on the implementation of the Lisbon Declaration on Youth Policies and Programmes, Catarina Camarinhas, Dwynette D. Eversley (LC/TS.2020/7, LC/CAR/TS.2020/1), 2020.
- 90. Proposal for a revitalized Caribbean Development and Cooperation Committee – Regional Coordinating Mechanism for Sustainable Development (CDCC-RCM): repositioning CDCC-RCM as the mechanism for sustainable development in Caribbean small island developing States (SIDS), Artie Dubrie, Omar Bello, Willard Phillips, Elizabeth Thorne, Dillon Alleyne (LC/TS.2020/6, LC/CAR/TS.2019/13), 2020.
- 89. Promoting debt sustainability to facilitate financing sustainable development in selected Caribbean countries: a scenario analysis of the ECLAC debt for climate adaptation swap initiative, Sheldon McLean, Hidenobu Tokuda, Nyasha Skerrette, Machel Pantin (LC/TS.2020/5, LC/CAR/TS.2019/12), 2020.
- 88. A preliminary review of policy responses to enhance SME access to trade financing in the Caribbean, Sheldon McLean, Don Charles (LC/TS.2020/4, LC/CAR/TS.2019/11), 2020.
- 87. Gender mainstreaming in national sustainable development planning in the Caribbean, Gabrielle Hosein, Tricia Basdeo-Gobin, Lydia Rosa Gén (LC/TS.2020/2, LC/CAR/TS.2019/10), 2020.

STUDIES AND PERSPECTIVES

Issues published:

96. Statistical disclosure control for Caribbean census tables
A proposal to expand the availability of disaggregated census data
Francis Jones
95. Economic implications of the ban on single-use plastics in the Caribbean
A case study of Trinidad and Tobago
Willard Phillips
Elizabeth Thorne
Camille Roopnarine
94. Creating an enabling environment for e-government and the protection of privacy rights in the Caribbean
A review of data protection legislation for alignment with the General Data Protection Regulation
Amelia Bleeker
93. The use of technology and innovative approaches in disaster and risk management
A characterization of Caribbean countries' experience
Luciana Fontes de Meira,
Omar Bello



Economic Commission for Latin America and the Caribbean (ECLAC)
Comisión Económica para América Latina y el Caribe (CEPAL)
www.eclac.org



LC/TS.2020/159