Proceedings of Statistics Canada Symposium 2018
Combine to Conquer: Innovations in the Use of Multiple Sources of Data

The Framework behind the Confidentiality Classification Tool and its Role in Modernization

Joseph Duggan, Jack Gambino, Claude Girard and Michelle Marquis[1]

Abstract

With the Confidentiality Classification Tool, Statistics Canada is implementing a small, but key supporting component of its recent Modernization Initiative. Sensitive statistical information (SSI), also referred to as confidential information, is now being classified along a continuum of risk, replacing the traditional binary classification that underpinned Statistics Canada's two separate working environments: one network for internal processing of protected information, and another network for external communication and the dissemination of its statistical products. The combination of this change along with other initiatives - relating to the use of alternative and combined data sources, modernized access to microdata, and bringing together more partners in collaborative efforts - will align our Disclosure Control practices with current and future IT infrastructures. The tool seeks to facilitate all of this by increasing awareness of confidentiality issues and practices, while helping data custodians determine a level of confidentiality for any selected data holding in Statistics Canada. This paper describes the methodology behind the intentionally-simple Confidentiality Classification Tool and the lessons that were learned in its development.

Key Words: Sensitive Statistical Information; Disclosure Control; Information Management; Microdata Access.

1. Introduction

The main body of this work describes the technical details of the tool, but it also includes information relating to the conceptual development of the framework for the tool. Following this introduction, and a few words on Modernization in Section 2, the tool is described in Section 3 and then specified in a more technical manner in Section 4. The paper concludes with some notes on the development of the tool and on its future.

1.1 A Note on Privacy and Confidentiality

Although related, the two separate concepts of privacy and of confidentiality merit clarification at this point. As described in Statistics Canada's Policy on Privacy and Confidentiality, the two are not equivalent. Privacy is the right to be left alone, to be free from interference, from surveillance and from intrusions. Governments have obligations with respect to the collection, use, disclosure, and retention of personal information. Confidentiality refers to a protection not to release identifiable information about an individual (such as a person, business or organization). This "trust" relationship between the supplier of the information and the organization collecting it is built on the assurance that the information will not be disclosed without the individual's permission nor without due legal authority.

Until very recently, confidentiality was generally perceived within Statistics Canada as a black and white issue, its binary nature was such that information was either confidential or it wasn't. This principle directed the way the organization worked, including how it structured itself and its information holdings.

[1]Joseph Duggan, Methodology Branch, Statistics Canada, Ottawa, Ontario, Canada, K1A 0T6; joseph.duggan@canada.ca; Jack Gambino, Statistics Canada jack.gambino@canada.ca; Claude Girard, Statistics Canada claude.girard@canada.ca; and Michelle Marquis, Office of Privacy Management and Information Coordination, Statistics Canada; michelle.marquis@canada.ca

2. Modernization at Statistics Canada

Statistics Canada's recent Modernization Initiative is based on five pillars: User-centric Service Delivery; Leading Edge Methods and Data Integration; Statistical Capacity Building and Leadership; Sharing and Collaboration; as well as a Modern Workforce and Flexible Workplace. Together, these pillars support Statistics Canada in fulfilling its mission statement: "Serving Canada with high-quality statistical information that matters." A more user-centric Statistics Canada will improve access, for statistical purposes, to more-relevant and more-timely data. This will be achieved by expanding modes of access and enhancing the discoverability of the data, so that the data are more available and the sources we have are more visible to researchers and other users. There will be a greater reliance on multiple and alternative data sources, along with new technology and new methodological approaches to using these sources and to the protection of confidentiality. These will be required to be in conformity with the commitment to leading edge methods and data integration. Key to these advancements will be the expansion of strategic partnerships with academic and policy researchers as well as with data providers and the continued building of statistical capacity amongst users. All can benefit from a better understanding of confidentiality, risk, quality, and the inherent advantages and limitations of the various forms of statistical information that Statistics Canada releases.

2.1 Increase in Access

The goal of increasing access to information must always be tempered by the requirement to protect the confidentiality of the data. At one extremity, complete access for everyone to all information would do nothing for confidentiality. At the other pole, complete confidentiality would only be assured by not allowing any access at all to anyone. In the course of modernizing, Statistics Canada has recognized the tension between these two potential states. Anil Arora, the Chief Statistician of Canada, stated that *"We are already learning that while we continue to change and modernize our work, our core values around confidentiality and high quality data remain the same* (Statistics Canada, 2018)*."* Statistics Canada seeks to increase access to information and still retain confidentiality as a core value by implementing the Confidentiality Classification Tool.

2.2 Classifying Confidentiality

The principle behind the Confidentiality Classification Tool is to define confidentiality along a continuum. Fundamentally, all confidential information remains confidential. However there would be degrees of confidentiality which would be mapped to degrees of protection and control of the data. As an input to the process, it was identified that Statistics Canada should take stock of what information it has. A lean and simple, self-administered tool was conceived to be used by statistical program practitioners, with the advice of experts on access and confidentiality, to complete the classification.

3. The Confidentiality Classification Tool

At its most basic level, the Confidentiality Classification Tool (hereafter, the tool) quantifies the level of disclosure risk for a statistical product and the sensitivity of the information it contains. Risk is evaluated across four commonly-studied types of disclosure (Duncan et al, 2011), slightly modified for the purposes of the tool. Sensitivity ratings were adapted from levels proposed from the Office of the Information and Privacy Commissioner of Ontario. These assessments of risk and sensitivity are combined to arrive at a score that determines the level of confidentiality for the product.

As set out in the introductory section of its User Guide, "The Confidentiality Classification Tool was developed to support Statistics Canada's modernization agenda… Sensitive statistical information (SSI) now needs to be classified along a continuum of risk. Replacing the traditional binary classification… will facilitate the implementation of the mobile workplace. It will also modernize access to microdata…." Additionally, one of the intended by-products of making the tool available to those engaged in the creation, storage, and use of data holdings within Statistics Canada is to bolster the awareness of confidentiality within the organization.

The ultimate objective was to develop a tool that would help determine a level of confidentiality for any selected data holding in Statistics Canada. A "data holding" was defined very broadly, it could include: a microdata file, table of statistics, statistical register, or anything that could potentially contain information that falls under the definition of

sensitive statistical information, such as a graph of data points or a draft codebook with frequencies that accompanies a survey data file.

3.1 Governance

The directors of divisions responsible for the information must sign-off on the review and attest to sufficient knowledge of risks amongst the reviewers in the making of theirs assessments. An internal microdata review and release committee will monitor the assessments from the metadata captured in the tool, and will compare results within programs and across the organization. Scores from applying the tool can be taken to the committee for appeal, or exception, depending on the circumstances and at the initiative of the data custodian.

3.2 Application

The tool was designed to be used by people who are knowledgeable about their data holdings but who are not necessarily experts in disclosure control methods. Programs could implement the tool for all of their existing products (disseminated and production files, information holdings), or could apply it on an as-needed basis. As examples, this ad hoc use could target new/redesigned programs and projects, special disseminations or collaboration/sharing initiatives, or even single files within their holdings. To facilitate the use of the tool, users are informed that it can be applied to a single product or to a group of products at the same time – following the structure that is encoded in the Generic Statistical Business Process Model (UNECE, 2013) as developed for the High-Level Group for the Modernisation of Statistical Production and Services in 2013.

3.3 Current form and content

At this point in time, the first version of the tool exists as an Excel workbook with four interrelated spreadsheets. The first sheet contains basic instructions and information selected from the separate user guide that accompanies the tool. The second sheet presents four questions to guide the assessment of disclosure risk and the third sheet is for the sensitivity assessment. The fourth and final sheet summarizes the information collected on the second and third sheets, and also presents the resulting overall score. The summary information includes a record of the particulars which identify the product being assessed and the principal assessor. Accompanying this information is the attestation by the director as described in section 3.1 above.

3.3.1 Disclosure risk questions

The sheet for evaluating disclosure risk does so using four questions following the same pattern, with one question for each of the four broad component types of disclosure: identify disclosure, attribute disclosure, inferential disclosure and residual disclosure, as set out below. The assessment of the risk from each disclosure type is to be made independently of the other broad types. In that evaluation they are guided to treat each type of disclosure as a "point of entry" into the data itself, and without considering all of the other safeguards in place to protect confidentiality.

The question for the *Identity Disclosure* type is prefaced by the description that it is at risk of occurring when confidential information can be revealed by means of the direct identifiers the product contains. The question reads: Can confidential information be revealed directly through the information displayed by the product? The assessment of risk is to be made in terms of the effort that is needed to arrive at a disclosure. The categories of response are: LOW risk as SIGNIFICANT effort is required; MEDIUM risk as SOME effort is required; and HIGH risk as LITTLE effort is required.

The tool describes *Attribute Disclosure* as associating data with a certain entity. This is typically done by grouping items described as indirect identifiers or "visible" information. The question that is posed by the tool is: Can confidential information be revealed by grouping attributes available in the product? And as done above, and for all the other broad types, the risk is evaluated in terms of effort using the same response categories.

Inferential Disclosure is said to arise when one can draw a certain conclusion with a high level of confidence that it applies to a certain entity. Can confidential information be revealed through a probabilistic statement derived from the information contained in the product? Typically, the products disseminated by statistical agencies are meant to support

inference, and as a result this type of disclosure is not typically as severe as the others. How this is reflected in the tool's scoring is described later in section 4.3.

Residual Disclosure takes into account the whole dissemination environment; assessed to the best of the assessors' knowledge as they are not expected to be aware of all information available publicly. Can confidential information be revealed by combining the information contained in this product with that from other products? It is noted that this definition of residual disclosure does not include notions of residual disclosure from within the product itself. Any instance of these 'internal' residual disclosures is an example of one of the forms given above and is to be assessed under the appropriate type.

3.3.2 Sensitivity question

In the section of the tool in which sensitivity is rated there is but one question to answer. Would the breach of the information in an identifiable format:

...cause severe harm to an individual or business?
...cause considerable harm to an individual or business?
...cause reputational damage or embarrassment to an individual or to a business?
...cause minimal harm to an individual or business?
...not cause any harm as the information is considered to be publicly available?

The selected response category must reflect an outcome that can reasonably be expected to occur, and not an extreme scenario. Note also that the impact on Statistics Canada of a breach was not taken into account here, nor elsewhere in the tool. This is not an omission, but rather an acknowledgement that this impact is judged to be significant regardless of the product being assessed with the tool.

4. Technical description of the tool

4.1 Design considerations

The different settings for the parameters used in the calculations can be modified in future versions of the tool. The underlying structure of the tool, and the model on which the calculations are based, were designed to be simple enough to be adapted in future versions of the tool, if required. Additionally, how this model is expressed can also be changed, as the way it is described here is just one of several possible ways.

4.2 Determining the classification level for confidential information

The classification level for a given data holding (or product), p, is a function of the Disclosure Risk score multiplied by the Sensitivity rating for that product. The first step simply presents the two factors (as in Table 4.3.1-1 below), and shows the result of their multiplication. In the second step, Table 4.3.1-2 acts as a look-up table for this result to arrive at the overall score (or classification level). A one-step look-up table would be simpler, but as the tool makes all of the calculations on its own, these tables are primarily used to demonstrate – for the assessors – the relationship between the two factors.

The Disclosure Risk score is the sum of component risks which are rated independently on four broad types of disclosure: identity disclosure, attribute disclosure, inferential disclosure, and residual disclosure. The ratings – either: low, medium, or high – are for the risk that each of these means of disclosure poses as a potential pathway to disclosure, under the assumption that the only controls against this are the ones already built into the data. The component ratings are transformed into values that depend on the type of disclosure under consideration. The values reflect the relative impact, and can be interpreted as weights. Their sum is capped at a maximum value of ten, for our purposes.

The Sensitivity rating is a straight transformation of the level of impact that a breach would have on an individual, business, or institution as a result of disclosure of information and is rated as either severe, high, medium, low, or negligible. These ratings assume values from five to one, in respective order.

4.3 Computing the classification level for confidential information

The classification level *C* for a data holding *p* is derived from the step function $F(DR_p \times S_p)$, where the Disclosure Risk score, $DR_p = \min [10, (Ident_p + Attr_p + Inf_p + Res_p)]$ is an element of $\{0,\ldots,10\}$ and the Sensitivity rating, S_p is an element of $\{1,\ldots,5\}$.

The four components of the Disclosure Risk vector ($Ident_p$, $Attr_p$, Inf_p , Res_p) are, respectively, the rating values for the risk of Identity, Attribute, Inferential, and Residual disclosure for p.

$$\text{Ident}_p = \begin{cases} 10, \text{if rated } \textit{High}, \\ 2, \text{if rated } \textit{Medium}, \\ 0, \text{if rated } \textit{Low}. \end{cases}$$

$$\text{Attr}_p = \begin{cases} 3, \text{if rated } \textit{High}, \\ 2, \text{if rated } \textit{Medium}, \\ 0, \text{if rated } \textit{Low}. \end{cases}$$

$$\text{Inf}_p = \begin{cases} 2, \text{if rated } \textit{High}, \\ 1, \text{if rated } \textit{Medium}, \\ 0, \text{if rated } \textit{Low}. \end{cases}$$

$$\text{Res}_p = \begin{cases} 3, \text{if rated } \textit{High}, \\ 1, \text{if rated } \textit{Medium}, \\ 0, \text{if rated } \textit{Low}. \end{cases}$$

The values assigned to each of the same three possible *descriptive* ratings vary to reflect the relative impact that each of these components was perceived to have on the total level of disclosure risk, as given by the Disclosure Risk score. These values can also be interpreted as weights.

The Sensitivity rating takes on the following values:

$$S_p = \begin{cases} 5, \text{if rated } \textit{Severe}, \\ 4, \text{if rated } \textit{High}, \\ 3, \text{if rated } \textit{Medium}, \\ 2, \text{if rated } \textit{Low}, \\ 1, \text{if rated } \textit{Negligible}. \end{cases}$$

And finally, the step function F can be evaluated using the product $(DR_p \times S_p)$ to obtain the resulting Confidentiality Level; where $C_p = F(DR_p \times S_p)$, as follows:

$$C_p = \begin{cases} 9, \text{if } DR_p \times S_p \in \{45, 50\}, \\ 8, \text{if } DR_p \times S_p \in \{36, 40\}, \\ 7, \text{if } DR_p \times S_p \in \{30, 32\}, \\ 6, \text{if } DR_p \times S_p \in \{21, 24, 25, 27\}, \\ 5, \text{if } DR_p \times S_p \in \{15, 16, 18, 20\}, \\ 4, \text{if } DR_p \times S_p \in \{10, 12, 14\}, \\ 3, \text{if } DR_p \times S_p \in \{7, 8, 9\}, \\ 2, \text{if } DR_p \times S_p \in \{4, 5, 6\}, \\ 1, \text{if } DR_p \times S_p \in \{1, 2, 3\}, \text{and} \\ 0, \text{if } DR_p \times S_p = 0. \end{cases}$$

4.3.1 Example computation

This subsection works its way through a mock example of the computation of the overall score.

If the DR vector of component risks (Identity, Attribute, Inferential, Residual) were assigned levels of (*low*, *medium*, *high*, *medium*) for each respective disclosure-type component, this would translate into the vector of values: (0, 2, 2, 1) which, when summed would give a total of 5 as the DR score.

Continuing the example, if the Sensitivity of the product was rated as *medium*, this would be converted to a value of 3. The product of the DR score and the Sensitivity rating would then be 15 (5 x 3) and the classification level for the data holding would be a 5 according to the step function displayed in Figure 4.3.1-1.

Figure 4.3.1-1
Step function relation between overall score and final classification level

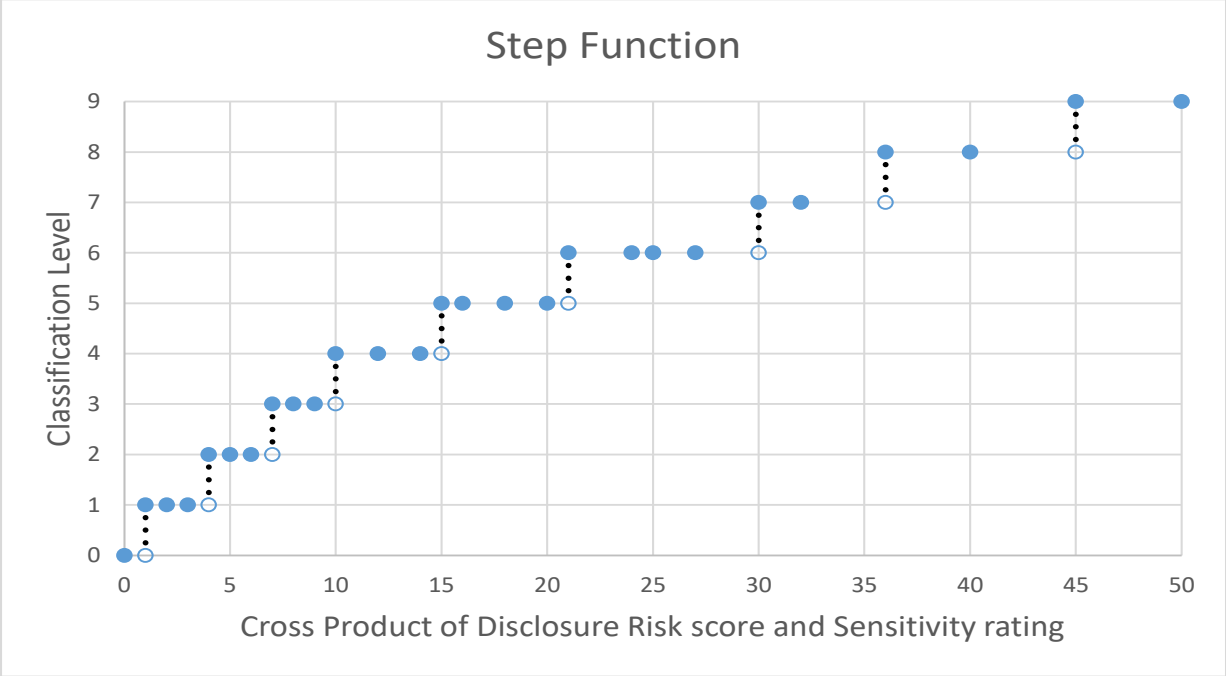

*The open circles and dashed lines typically used only for continuous-valued step functions are shown above solely to make the steps more apparent.

As an aside, it is an interesting interpretation that the Disclosure Risk score can be pictured as a maximal length within a multi-dimensional space created by potential risks. Note that if the values used for the ratings given to the four components of the Disclosure Risk vector ($Ident_p$, $Attr_p$, Inf_p , Res_p) are instead replaced by their square roots, then the Disclosure Risk score would be equivalent to the square of the Euclidean length of the disclosure risk vector – capped at 10. In that case, for the example above, ($Ident_p$, $Attr_p$, Inf_p , Res_p) would equal $\left(0, \sqrt{2}, \sqrt{2}, 1\right)$. The length of this vector, or the square root of its DR score, is then $\left(\sqrt{0}^2 + \sqrt{2}^2 + \sqrt{2}^2 + \sqrt{1}^2\right)^{\frac{1}{2}}$, or $\sqrt{5}$. Multiplying this distance by the Sensitivity rating defines a volume for the complete confidentiality level.

Rather than the step function, the tool presents the overall score in an easier to understand tabular format. This table format uses colour in a gradient from green at the low risk end of the spectrum to red at the high risk area diagonally opposite at the upper right-hand corner of the table. The values found inside the table naturally result in a table that is more green than red. The pre-conditioned associations with these colours are meant to re-inforce confidence in the lower risk evaluations and well-placed concern in regards to the higher risk area. The relative sizes of these areas are meant to reflect the Modernization objective of increased access under recognition that confidentiality has been assessed properly for all products.

Table 4.3.1-1
Determination of overall score from disclosure risk score and sensitivity rating

Sensitivity	**5**	0	5	10	15	20	25	30	35	40	45	50
	4	0	4	8	12	16	20	24	28	32	36	40
	3	0	3	6	9	12	15	18	21	24	27	30
	2	0	2	4	6	8	10	12	14	16	18	20
	1	0	1	2	3	4	5	6	7	8	9	10
		0	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
		Disclosure Risk										

Table 4.3.1-2
Tabular determination of classification level based on overall risk-sensitivity score

Final Classification Levels based on Combined Disclosure Risk and Sensitivity Factors:		
Lower Bound	Upper Bound	Level
45	50	9
36	44	8
30	35	7
21	29	6
15	20	5
10	14	4
7	9	3
4	6	2
1	3	1
0	0	0

5. What is past, is prologue

5.1 Retrospective

By means of summary, it may be instructive to review a few of the major milestones in the development of the tool. Early attempts to create a "checklist" to guide the assessment of confidentiality levels spawned classification-tree style solutions. These trees quickly grew large and unwieldy as allowances were made for different types of data (social/enterprise/institutional, survey/administrative data, etc.).

As an alternative, it was proposed to use the four broad types of disclosure as a framework, and to put aside considerations that were specific to different forms of data. When wrestling with the question of how to relate these separate measures into a Disclosure Risk score, the application of weights was the logical, methodological direction to take – but required leadership and consultation to arrive at the function and parameters documented in this paper. Setting the concept of sensitivity aside until it was agreed that disclosure risk had been settled, proved almost providential, given that the consideration of the sensitivity/value of something seems to be ingrained in how human nature assesses risk. As dimensions of confidentiality, it also seemed natural to use a two-dimensional grid to combine the factors and present their relationship. The last milestone worthy of mention was another generalization: the use of the General Statistical Business Process Model to create similar classes of data products from a program area and reduce the number of assessments required in practice.

5.2 Future developments

To conclude, the reminder is given that confidentiality is, as always, a core value at Statistics Canada. The Confidentiality Classification Tool is meant to be, first and foremost, a simple tool. There are plans to make the tool web-based in the future and to have the results of the assessments available as metadata included in the statistical products' documentation. The tool is also meant to expand awareness of how confidentiality is maintained while access to statistical information is increased. It is only one little piece of the larger picture of Modernization within Statistics Canada, and it will – and has been – shared with other departments and agencies. Ultimately, and as intended all along, it will be adapted and improved over time.

Acknowledgements

The authors would like to recognize the participation of the other members of the working group, and its management, who collaborated in the development of the Confidentiality Classification Tool. Additionally, several program areas need to be thanked for their work in running selected data products through test runs of the tool and user guide. Lastly, François Brisebois needs to be thanked for his insightful contributions to the project and for his helpful reviews of this paper.

References

Duncan, G., Elliot, M. and Salazar-Gonzalez, J. (2011), *Statistical Confidentiality: Principles and Practice*, New York: Springer-Verlag.

Federal Committee on Statistical Methodology (2005). "Statistical Policy Working Paper 22 (Second version): Report on Statistical Disclosure Limitation Methodology." NTIS PB94-165305.

OECD Glossary of Statistical Terms. https://stats.oecd.org/glossary/index.htm

Statistics Canada (2018), Modernization Bulletin: June 2018, internal monthly newsletter, Ottawa, Canada

Statistics Canada (2018b), "Confidentiality Classification Tool: User Guide (version 1.0 external)", Ottawa, Canada

UNECE (2013). Version 5.0 of the Generic Statistical Business Process Model (GSBPM) https://statswiki.unece.org/display/GSBPM