Government Statistical Service

Best Practice and Impact

GSS Methodology Advisory Committee

# Privacy and Data Confidentiality Methods

## National Statistician's Quality Review

Keith Spicer
Gentiana D. Roarson

**December 2018**

## Foreword

The rapid increase in the detail, volume and frequency of data collected, alongside the diversification of data sources available, presents a real opportunity for the statistical community to innovate. This richer detail will provide better statistics that deepen our understanding of society and better support decision making.

On the other hand, making more data available raises new concerns and new challenges to protecting privacy and confidentiality of personal information. In this challenging and rapidly changing landscape, the statistical community has a legal and an ethical obligation to protect the confidentiality of data, while at the same time striving to meet evolving user demands for more detailed and helpful statistics.

I welcome the findings of this National Statistician's Quality Review (NSQR) of Privacy and Data Confidentiality Methods, which will help the Government Statistical Service (GSS) to take full advantage of the cutting-edge developments and research conducted by world leading experts from across academia and the private sector. Closely collaborating and joining forces with these experts helps us to ensure that the GSS prepares for the future and identifies opportunities to improve and innovate.

## Joining forces with leading experts

We are delighted to collaborate with leading experts in privacy and confidentiality external to the GSS and grateful for their contribution: Professor Natalie Shlomo and Professor Mark Elliot from the University of Manchester, Professor Josep Domingo-Ferrer from Rovira I Virgili University in Tarragona Spain, Professor Kerina Jones and Professor David Ford from Swansea University, Professor Jim Smith and Professor Felix Ritchie from the University of the West of England, Dr Hector Page and Charlie Cabot from Privitar and Professor Kobbi Nissim from Georgetown University, USA.

## The case for change

The Digital Economy Act (DEA) 2017 enables better sharing and use of data across organisational boundaries, at a time of dramatic increase in the volume of data available and a wide range of sources. New and updated legislation, including the General Data Protection Regulation (GDPR), have brought about major changes to the way organisations process personal data, encouraging greater transparency and accountability.

These developments present an unprecedented opportunity to innovate with data, while safeguarding privacy and fostering public trust. On the other hand, against a backdrop of

increasingly sophisticated attacks being developed by intruders and the potentially serious consequences of data breaches, the data revolution presents significant new challenges to protecting privacy and confidentiality.

It is vital that the statistical community understands and addresses these evolving challenges to provide a solid foundation for innovation to take place. This is not a straightforward task, and there is a need to draw upon the full range of expertise in this fast-developing field.

By joining forces with leading experts, this National Statistician's Quality Review explores the latest advances in these methods. From a suite of contributing articles setting out the latest developments, the review draws out emerging themes and articulates the challenges facing the statistical community. It also sets out the steps that need to be taken across the statistical system, which will be further developed and implemented through engagement across organisations.

## Challenges faced by producers of statistics

In collaboration with leading experts, this review has explored the latest research and trends for privacy and confidentiality methods and has identified the following challenges currently faced by producers of statistics:

1. Assessing and balancing what is theoretically possible with what is feasible and appropriate in a statistical context. In practice, balancing risk of disclosure with data utility is a key consideration.
2. Predicting and planning for future change when the nature of that change is uncertain and public attitudes to privacy are changing and evolving; this is heavily influenced by publicity around any data privacy breaches.
3. Evaluating the risk of intruder attacks and new types of privacy attacks, particularly from linking with the information available from other sources.
4. Keeping up-to-date with the latest methodological research and technological advancements, and building capability.
5. Identifying and developing approaches and methods that better exploit the potential of new data sources and technological advancements to investigate the feasibility of new developments whose practical usefulness has not yet been clearly demonstrated (such as machine learning and artificial intelligence applications, or practical applications of differential privacy).
6. Assessing and developing the potential of specialist software and automation.
7. Communicating disclosure risk, choice of privacy and confidentiality methods and their trade-offs to the statistical community, decision makers and the public.
8. Future-proofing data releases and considerations around what is going to (or likely to) be released in the future by different data providers. Legally, producers of statistics do not have to consider any data not already (or due to be) in the public domain, but best practice should be to consider future-proofing in expectation of what is likely to be released.

This review aims to support decision makers as well as privacy and data confidentiality experts from across the Government Statistical Service (GSS) to understand the art of the possible for these methods and make informed decisions on the next steps required to keep pace with latest developments.

Follow-up work will help to expand understanding of how these methods can be developed and implemented across the GSS. Further strands of work are to be confirmed, but it is expected that they will support activities such as identification of priority areas, further understanding of the current practices and development of jointly agreed implementation plans.

## Proposed GSS next steps

The greatest challenge for all specialists in confidentiality protection and for producers of statistics is to move with the times. An important outcome from this review must be to help make inform decisions on further developments needed for these methods and to build capability across the Government Statistical Service (GSS). The key strands are:

1. Establish a disclosure control centre of expertise led by Office for National Statistics (ONS). The centre should look to build expertise across the GSS and should act as a central point of contact for disclosure control advice, including reviews of the current methods used against best practice and developments in this field
2. Further expand collaborations between the GSS, academia, private sector and other National Statistics Institutes (NSIs). This will include setting up an international working group with other NSIs to share knowledge, research and future plans on privacy and confidentiality. Engagement and collaboration with experts in academia will also be utilised to ensure the most up to date methods are understood, researched and applied across the GSS.
3. Reinforce and encourage a balanced disclosure risk vs. data utility approach. Develop tools and capability for risk assessment, including approaches for addressing and evaluating the risks brought by new types of intruder attacks. Ensure there is good practice in risk management rather than only focusing on risk reduction. Acknowledge that the risk can never be zero and have procedures in place in the event of a confidentiality breach, including a more formal breach policy.
4. Expand the tool kit of courses and guidance materials, ensuring they meet user needs. This could include provision of e-learning and other activities such as clinics and workshops, drawing on existing and emerging expertise across GSS. One of the first tasks here will be to assess what training and guidance are currently used.
5. Explore the applications of new techniques and technology to identify and develop new approaches and methods that take better advantage of the opportunities provided by the data environment. Identify exploratory research needed to demonstrate the practical application for these new approaches for the production of statistics and highlight their trade-offs, advantages and limitations.
6. Explore the potential for practical applications of new software and technologies in protecting privacy and confidentiality. Coordinate the development of specialist

software and automation to support the assessment and implementation of anonymisation best practice. Explore options for applying SDC methods to support the development of new dissemination strategies, such as automated table generators, determining their applications and limitations.

7. Develop an overarching confidentiality protection framework that offers data providers a way of measuring the risk and data utility across SDC methods and privacy models in a more unified way. Explore options for this confidentiality framework to provide anonymisation transparency and help communicate disclosure risk, choice of privacy and confidentiality methods and their trade-offs to the decision makers and the public.
8. Establish a GSS Task Force to implement the next steps from this NSQR. This forward-focused group should report to GSS Statistical Policy and Standards Committee, and draw on the expertise and experience in anonymisation practice from across the GSS.

## Further research and development

There is a need to continue to build on the research already carried out as part of this NSQR, working across the Government Statistical Service (GSS) and engaging support from the academic community and private sector. Suggested topics for the forward programme of research include:

- Investigate how statistical disclosure control methods may be adapted to learn from insights provided by differential privacy by conducting a pilot on a “non-sensitive dataset”.
- Assess whether disclosure risk is increased by providing the parameters of the protection method.
- Review the set of intruder scenarios to include potential and emerging threats to privacy and confidentiality.
- Develop strategies for defending against reconstruction attacks.
- Assess the risk pertaining to inferential disclosure.
- Analyse the threat of social media and other data sources to identify appropriate strategies to protect data releases against the risk of a linking attack.
- Assess the practicality of producing and using synthetic datasets.
- Further develop the measures to quantify risk of disclosure and information loss including taking context into account.
- Explore the practical applications of machine learning and artificial intelligence protection of privacy and confidentiality and options for avoiding the replication of human biases.

## Summary of GSS challenges, next steps and further research and development

| | GSS challenge | GSS next steps | GSS further research and development |
|---|---|---|---|
| 1 | Balancing theory with what is possible/feasible in a statistical context | Disclosure control centre of expertise led by the Office for National Statistics (ONS) | Develop the measures to quantify the risk of disclosure and information loss including taking context into account |
| 2 | Predicting and planning for future change | Expand collaborations between the Government Statistical Service (GSS), academia, private sector and other National Statistical Institutes | |
| 3 | Evaluating the risk of intruder attacks | Develop tools and the capability for risk assessment | Review the set of intruder scenarios to include potential and emerging threats to privacy and confidentiality<br><br>Develop strategies for defending against reconstruction attacks<br><br>Assess the risk pertaining to inferential disclosure<br><br>Analyse the threat of social media and other data sources to identify appropriate strategies to protect data releases against the risk of a linking attack |
| 4 | Keeping updated with the latest methodology | Expand the tool kit of courses and guidance materials | |

| | | | |
|---|---|---|---|
| 5 | Identify and develop new approaches/methods that take better advantage of the opportunities provided by the data environment | Identify exploratory research needed to demonstrate the practical application for the approaches/ methods | Investigate how statistical disclosure control (SDC) methods may be adapted to learn from insights provided by differential privacy (DP) by conducting a pilot on a “non-sensitive dataset”<br><br>Assess the practicality of producing and using synthetic datasets<br><br>Explore the practical applications of machine learning and artificial intelligence protection on privacy and confidentiality |
| 6 | Assessing and developing the potential of specialist software | Develop specialist software and automation to support the assessment and implementation of anonymisation best practice | Develop the measures to quantify risk of disclosure and information loss including taking context into account |
| 7 | Communicating disclosure risk | Develop an overarching confidentiality protection framework | Assess whether disclosure risk is increased by providing parameters of the protection method<br><br>Further develop measures to quantify risk of disclosure and information loss including taking context into account |
| 8 | Future-proofing data releases | Ongoing work including Statistical Standards and Policy Committee (SPSC) task force | Analyse the threat of social media and other data sources to identify appropriate strategies to protect data releases against the risk of a linking attack |

## Latest research and emerging themes

### Rapid developments in data collection and data processing bring new challenges

The rapid growth in the number and range of bodies releasing data has made it increasingly difficult for the Government Statistical Service (GSS) to keep track of relevant publicly available data.

The push for open data and transparency has been countered by the increase in privacy and confidentiality concerns and the increase in threat, both perceived and real, posed by multiple data sources from multiple data providers and possible linkage between those sources. The threat has been exacerbated particularly by the growth in use of social media and people's willingness to post considerable detail about themselves on these platforms. It is possible that data collected and published by a government department could be matched to information posted on social media, leading to confidential details being revealed. Elliot and Domingo-Ferrer discuss that in practice it is difficult to determine which matches are incorrect so the assumption is made that any possible matches constitute some level of disclosure risk.

### Linking data can increase the risk of disclosure

The Digital Economy Act (DEA) and developments in data technology will make it easier for the Government Statistical Service (GSS) to harness the power of data by linking and matching two or more data sets. Linking data sets can be a methodologically challenging, multi-stage process and it is important that privacy and confidentiality are protected at all stages.

Jones and Ford discuss the different linking methods with their advantages and disadvantages, while Ritchie and Smith discuss the potential of re-identification through linking and unpredictability of future scenarios. The choice of data linking method depends on a variety of factors, including burden on the data provider and technological limitations. Some of the measures for managing disclosure risk are specific to data linkage.

Identifiable data are required at some stage in all data linking methods and the resulting linked data sets are not immune to disclosure risks. As intruder attacks are becoming more sophisticated and easier to conduct with the aid of new technologies, it is essential to identify vulnerabilities at each stage of the data linking process and to develop approaches to protecting the data that will minimise the risks to an acceptable level.

There is promising initial evidence for new solutions to issues resulting from linking diverse datasets. Ritchie and Smith discuss the advantages and challenges posed by using new techniques such as machine learning based record linking, highlighting that the same methods and analytical tools can also be used in malicious linking attacks. All

these techniques are in early stages of development and will require more research before providing sufficient evidence that they can become viable options.

## Balancing risk of disclosure with data utility

In achieving high data utility, it is important to maintain a balance of risk and utility in outputs and avoid being too risk averse. Defining what this balance should look like in practice is not straightforward, and guidance on this is not clear.

The assessment is almost always dependent on the risk appetite of the data provider or the Information Asset Owner (IAO). It is important to be able to provide good advice and guidance to IAOs to ensure practice that is, wherever possible, consistent across the Government Statistical Service (GSS) and encourages risk management as opposed to complete risk avoidance.

Statistical disclosure control (SDC) methods should minimize the risk of disclosure to an acceptable level (in line with legal requirements and good practice) while releasing as much information as possible. The decision of which SDC methods to use is dependent on trade-offs between minimising the disclosure risk and having the least possible adverse impact on the usefulness of data.

This decision also needs to consider the different types of data and contexts (referred to as the data environment), including previously released information and intruders' knowledge of the population. Elliot and Domingo-Ferrer discuss the relationship between the data environment and the data intended to be released, and draw attention to the inherent difficulty of assessing risk of disclosure given that the data held in private databases are a key source of uncertainty for the SDC process. Shlomo emphasises that to achieve the optimal balance between risk and utility it is important to have quantitative measures for both risk of disclosure and information loss.

The impact of SDC methods used on the utility of the data can be communicated through a disclosure risk – data utility map. SDC application entails using different methods with different parameters/thresholds and the resulting disclosure risk and data utility should be quantified, to support the decision on the best risk-utility reconciliation.

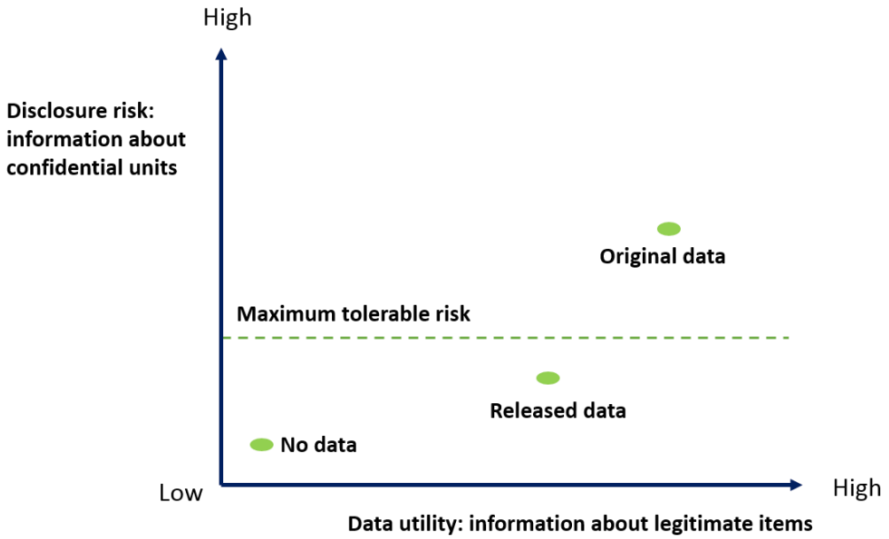


**Figure 1: The statistical disclosure control problem**

Elliot and Domingo-Ferrer discuss privacy models and their application to control for disclosure risk. The privacy models (for example k-anonymity) can be seen as alternative approaches to SDC methods, but in practice they can be strengthened by using one or several SDC methods.

One of the latest additions to privacy models is differential privacy (DP). Page, Cabot and Nissim define DP as a privacy model that bounds privacy risk and formalises the intuitive view about privacy that a statistical output should reveal (almost) no information specific to an individual within the data set.

More powerful intruder attacks including new types of attack (such as reconstruction attacks) are some of the motivating factors that led to the development of DP. Protecting against disclosure risk using differential privacy is a much more context-free approach and is easily quantifiable once the parameter in the model is set, but setting this parameter is one of DP's practical challenges.

As DP is a rather new addition to privacy models, it is still in the early stages of transition from research to practical applications for statistical production, and users are still learning how to apply it effectively in practice. See Page, Cabot and Nissim and Shlomo.

Elliot and Domingo-Ferrer point out that these privacy models should not be regarded as competing models and an understanding of their applications and limitations can inform the use of one model to boost another.

An added difficulty in developing a unified framework that offers data providers a way of measuring the risk and data utility across SDC methods and privacy models is inconsistency in their application in practice; in part a result of varying risk appetite, with individual context being part of the risk assessment. Whereas data with 'zero risk' (none

or negligible) or 'total risk' (identified) are relatively easy to define, something in-between which provides sufficient utility while still providing sufficient risk protection is very much dependent on the relative importance subjectively put on the two requirements.

Elliot and Domingo-Ferrer conclude that, in general, methods developed to quantify and understand data utility are less developed than methods designed to measure disclosure risk. Overall more research is needed for both types of methods to assess whether they are satisfactory. Further innovation and development will better support decision making and provide solutions to the tension between demand for more detailed information and the obligation to ensure privacy and confidentiality. This becomes more pressing when considering big data, data mining and a rapidly expanding data environment.

## Intruder testing is a potential approach to managing threat

Across the Government Statistical Service (GSS), statistical disclosure control (SDC) methods are selected to be consistent, practical and implementable in a reasonable timeframe with available resources. Shlomo provides a comprehensive description and further examples for the different SDC approaches.

Traditionally, a statistician's choice of SDC method (or combination of methods) and threshold parameters is guided by expert advice and previous good practice examples of dealing with disclosure risk and information loss for similar statistics. For example, for frequency tables, record swapping is the preferred pre-tabulation method (implemented on microdata prior to constructing the table) due to ease of implementation, but table redesign, random rounding or noise addition is advised for post-tabulation treatment (after the table is constructed). For magnitude tables, primary and secondary suppression is advised with consideration given to the sensitivity of the disclosive cell and threshold rules (see Shlomo for a classification of disclosure risks). A high-level classification of these methods by type of output can also be found in Elliot and Domingo-Ferrer.

Intruder testing can help identify vulnerabilities and help measure disclosure risks. However, this is a complex and resource intensive process. Elliot and Domingo-Ferrer discuss the different stages, expertise and resource required, and recommend this approach for new data situations where the calibration of disclosure risk measures is difficult to achieve.

A disclosure risk that is becoming more prominent with the development of automatic table generators and remote analysis servers is inferential disclosure. More widely, the inferential disclosure risk has led the statistical community to recognise the increased need to use perturbative methods (see Shlomo and Page, Cabot and Nissim). Inferential disclosure risk is an area where further research is needed to explore options and identify solutions that can help inform the choice of methodology used to protect the different types of statistics. Overall, further research is needed to help develop approaches that can deal with, and build protection against, malicious behaviours and intruder attacks.

Elliot and Domingo-Ferrer discuss the challenges posed to the existing SDC model by the changes in the data environment, and the increase in the data available. Contributing to this challenge is the increase in capacity to process data (such as linking). However, minimising the risk of disclosure when working with different types of data sets has exposed some of the limitations of the standard SDC model and the need to develop new approaches and identify new risks.

Machine learning could be a useful tool to help further develop the SDC model but this is at an early stage of development.

As SDC methods are being developed and evaluated, threat models are updated to reflect what are considered as potential threats at a particular point in time. These models have to be regularly reviewed to allow for new developments, especially increase in computing power and richness of auxiliary data.

## The potential of synthetic data

Synthetic data have the properties of real data sets but are composed of artificial individual cases generated using a particular model. Elliot and Domingo-Ferrer touch on the lack of consensus between experts on whether synthetic data can be viewed as a statistical disclosure control (SDC) method or as an alternative to SDC. In practice, a synthetic dataset is as good as the model underpinning it; it is difficult to capture conditional relationships between variables and the risk of disclosure is not completely eliminated. Elliot and Domingo-Ferrer discuss the potential of machine learning in synthetic data generation and model building, and the initial research conducted to synthesise data using deep learning techniques.

Page, Cabot and Nissim identify differentially private synthetic data as one path for the differential privacy model deployment, and discuss some potential practical advantages such as allowance for indefinite queries of the data.

## Inconsistent terminology and guidance can be a source of confusion

It is relatively common for experts and researchers to use different terms to refer to the same concept, or alternatively, to use the same term when referring to different concepts. While some of these terms have legal implications, they are not always used consistently. To complicate matters further, in this fast-evolving field there isn't always consensus between experts on how some of the new developments can be implemented in practice. In addition, the terminology used is evolving.

Anonymisation, privacy, confidentiality and statistical disclosure control (SDC) are sometimes used interchangeably. SDC methods are just one of the tools in the anonymiser's toolbox, and to fully identify privacy concerns an assessment that goes beyond disclosure control is needed.

Elliot and Domingo-Ferrer differentiate between anonymisation (transforming personal into non-personal/anonymised data) and SDC methods, with the latter being a subset of methods used as part of the anonymisation process to manipulate and minimise disclosure risk.

This makes the dissemination of findings to a non-expert audience more difficult, and is at times a source of confusion that could be avoided by developing guidance and encouraging the consistent and clear use of terminology across the Government Statistical Service (GSS) and wider. When communicating to the public how risks of disclosure are dealt with, consistent and meaningful language is essential in building public trust and alleviating concerns around privacy breaches.

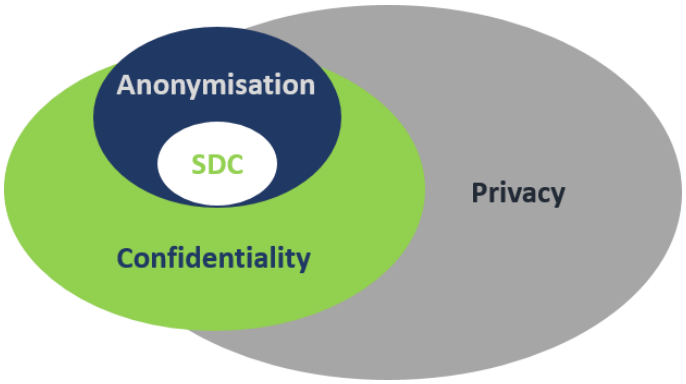


**Figure 2: The relationship between privacy, confidentiality, anonymisation and SDC**

(Credit: Mark Elliot and Josep Domingo-Ferrer)

## New dissemination strategies

Increasing demand for more accessible and detailed statistical data has led to the development of new methods for dealing with privacy and data confidentiality threats. This work resulted in new and innovative dissemination strategies: remote access, data enclaves and web based dissemination tools such as flexible/automated table generators, and remote analysis servers**.**

Shlomo discusses how flexible table generators development is driven by demand from policy makers and researchers. The tables allow them to define their own outputs from a set of pre-defined variables and categories. The automatically generated table is checked against a list of criteria and if these criteria are met it is released to the researcher without the need for human intervention.

Remote analysis servers are online systems that provide similar outputs to flexible table generators, without the need for human intervention. With remote analysis servers all statistical disclosure control (SDC) methods are applied pre- or post tabulation based on the rules and thresholds programmed in the system (see Shlomo and Ritchie and Smith).

The challenges for the future development of SDC methods are to examine the potential of privacy guarantees pertaining to differential privacy and to develop applications for new and innovative statistical dissemination strategies.