



Introduction

to Anonymisation

Bristows



Contents page

| | |
|---|-----------|
| 1. Introduction | 3 |
| 2. Executive Summary | 5 |
| 3. De-identification techniques | 8 |
| 3.1. Masking, linkage attacks, and k-anonymity | 9 |
| 3.2. k-anonymity, Big Data, and environmental controls | 10 |
| 3.3. Aggregates, reconstruction attacks and differential privacy | 12 |
| 3.4. Machine learning models and model inversion attacks | 14 |
| 3.5. Synthetic data | 15 |
| 3.6. Balancing competing priorities | 16 |
| 4. Anonymisation: the legal position in the UK | 18 |
| 4.1. Legislation | 18 |
| 4.2. Guidance from regulators | 21 |
| 4.3. Case Law | 22 |
| 4.4. When is data anonymous? | 23 |
| 5. Summary of issues with the UK approach | 26 |
| 5.1. The definition of personal data | 26 |
| 5.2. Is there a realistic possibility of an attacker? | 28 |
| 5.3. How to account for the data environment when assessing re-identification risk? | 30 |
| 5.4. According to the guidance, when is a re-identification attack considered successful? | 31 |
| 6. Divergence between the UK and EU positions | 33 |
| 6.1. Requirement to delete the source data in order to create an anonymous extract | 34 |
| 6.2. Identifiability: the regulators' perspective | 34 |
| 6.3. The threshold for identifiability; the regulators' perspective | 35 |
| 7. Conclusion - What does this mean for organisations? | 37 |
| Annex A - Relevant case law | 40 |
| Annex B - Common De-identification Techniques | 46 |
| Data Transformations | 46 |
| Environmental Controls | 47 |
| About | 48 |
| About the authors | 49 |

1. Introduction

Data protection law applies to personal data, which the GDPR defines as “any information concerning or relating to an identified or identifiable natural person”.¹ In contrast, anonymous information, in other words information that can no longer be connected to an individual, is outside of the scope of data protection law. Anonymisation therefore defines the perimeter of data protection law, with significant consequences for organisations.



Organisations are becoming more data driven. Evolving technology and proliferation of data means that more of the data that an organisation uses could be considered ‘identifiable’ within the meaning of data protection law. A broader spectrum of professionals are using or making decisions about data and therefore needing to interact with the data protection regime. Projects using data on a large scale can deliver significant benefits for the organisation, but carry a correspondingly high level of data protection risk if the data is considered identifiable. And many projects involve international collaboration, requiring consideration of many different, sometimes inconsistent, data protection laws.

In our experience, organisations often consider anonymisation as they seek to use more of the data they hold while managing compliance obligations. This raises several practical questions, including: how to anonymise data, and how to know whether data is anonymous? The first relates to the technical tools at an organisation's disposal. The second is a legal question about the point at which data is anonymous. These questions can be challenging, making anonymisation a source of confusion for many organisations. Our paper is a summary of the most important things we have learned in our combined experience advising on these challenges.

The law and guidance on data protection issues is evolving. The law is untested in many areas. The Information Commissioner's Office (ICO) and European Data Protection Board (EDPB) have both announced forthcoming guidance on anonymisation. The current ICO guidance, published in 2012, predates GDPR. We see a significant opportunity for regulators to clarify the situation, and look forward to engaging with them as new guidance is developed. The courts are also driving change. Cases continue to emerge on the parameters of ‘anonymised’ versus personal data. However UK case law is not always consistent, is at odds with some of the guidance, and mostly predates the GDPR.

¹ The European Union (Withdrawal) Act 2018 transposed the GDPR into UK law as the UK GDPR. The UK GDPR has undergone minor changes to enable it to operate as an independent piece of UK law (e.g. removing references to ‘the Union’), but there are currently no substantive differences between the UK GDPR and the GDPR. We use ‘GDPR’ in this paper to mean the EU GDPR and UK GDPR, unless otherwise stated.

This paper is an introduction to anonymisation for anyone responsible for using data that may be 'identifiable', whether in a commercial, scientific research or any other setting. It is also aimed at those responsible for ensuring compliance with data protection rules, such as privacy, data protection, legal, compliance, risk, or information governance professionals. It analyses the UK's legal and regulatory approach to anonymisation, and the key technical and business issues organisations seeking to anonymise data should consider.²

² For ease of reference, throughout we refer to "organisations". Generally in doing so we are referring to the parties responsible for making the decision as to the purpose for which and manner in which data is used, that is, data controllers, while acknowledging that processors may also be involved in the issues discussed.

2. Executive Summary

Anonymisation is sometimes presented as a silver bullet; a way to use data outside of the scope of data protection law. However, in practice anonymisation is complex and is not always appropriate. This paper summarises the relevant technical controls and legal considerations organisations wanting to anonymise data in the UK should be aware of, and highlights some current areas of uncertainty.



Anonymisation involves applying controls to make it less likely that an individual could be identified in a dataset, and then assessing whether the risk of re-identification is sufficiently low to cross the legal threshold for the data to be considered 'anonymous'. Controls can act on the data itself (e.g. removing identifying attributes or blurring data to make it less accurate) or on the environment in which the data is used (e.g. allowing only vetted employees to access the data in a secure location).

We will show that anonymisation can be difficult to achieve because the technical methods for re-identifying data (which we call re-identification attacks) are evolving (Section 3) and because the legal test for whether data is anonymous is risk-based, leaving room for interpretation (Sections 4 and 5).

Section 3 explores these controls in detail, using masking, *k*-anonymity, differential privacy, privacy-preserving machine learning and synthetic data as examples. We discuss the scenarios in which each control could be relevant. We show that applying controls only to the data itself will often lead to a significant utility and usability trade off.

The utility and usability trade off means that it may not always be possible to anonymise data while retaining sufficient utility for the intended purpose. Achieving anonymisation through a combination of controls that act on the data itself and controls on the environment in which the data is processed can allow organisations to find the trade off with utility and usability that works best for them.

Section 4 considers anonymisation as a legal concept, introducing the relevant legislation, case law, and guidance. The UK test for anonymisation features, amongst others, the 'motivated intruder' test. In other words, could a reasonably competent 'motivated intruder' successfully re-identify an individual in the de-identified data?

Section 5 highlights specific challenges for organisations navigating the UK's legal position. These include inconsistency in the case law, divergent definitions of personal data between the 1998 UK Data Protection Act (on which most existing UK case law and guidance is based) and the GDPR, and aspects of the motivated intruder test.

Section 6 considers areas of legal divergence between the UK and EU positions. Such divergence is a significant challenge for organisations operating across multiple jurisdictions. Regulators in different jurisdictions may use different interpretations of 'anonymous' and could reach different conclusions on whether data is anonymous in a specific scenario.

We also show that anonymisation may not be necessary or indeed may make it harder to achieve some commercial or operational goals. We show that where it may be necessary or desirable to use identifiable data, data protection need not be an insuperable obstacle. The techniques described in Section 3 can be used to anonymise data. They can also be used as safeguards to reduce identifiability.

Applying safeguards, such as pseudonymisation, may help open up more use cases while remaining within the data protection regime. For example, the legitimate interest lawful basis includes a balancing test, where safeguards play a role in determining whether the proposed use of data is permissible.

Section 7 concludes with recommendations for organisations considering anonymisation. In many cases, organisations seeking to use data they hold for a secondary purpose may have several options, including anonymisation. The GDPR explicitly recognises these and provides a legal foundation for them in the field of scientific and other types of research. Accordingly, in the context of such research, the introduction by organisations of safeguards that restrict the ability to identify an individual is recognised as an important factor in limiting the circumstances in which, from a legal perspective, data is considered identifiable.

These options will have different implications for utility, usability, and ongoing legal obligations. As such there is no one strategy for all datasets and use cases. Organisations should start by defining their requirements, identifying the range of possible options, assessing the trade-offs involved with each and selecting the most appropriate for the specific scenario.

Busting four myths about anonymisation

1. “Anonymisation is easy and always an option”.

Anonymisation is sometimes presented as a silver bullet; anonymise your data and then you can do anything you want with it.

The reality? Sadly, it is not. It is a useful legal option, but will only be the right answer some of the time. Effectively de-identifying data, while maintaining sufficient utility and usability for your use case, and accurately assessing re-identification risk, is hard, easy to get wrong, and not always possible.

2. “Pseudonymisation = anonymisation”.

We often hear the terms used interchangeably or imprecisely. For example, saying that removing direct identifiers is the same as anonymising data.

The reality? Pseudonymisation means processing data such that the data can no longer be attributed to a specific individual without the use of additional information. In practice, this often means removing direct identifiers from the data. The examples of re-identification attacks in Section 3 show that it will often be trivial to re-identify data that has been pseudonymised and released. In some cases, protecting pseudonymised data by placing it into a controlled environment can reduce re-identification risk to the extent that the data is considered anonymous. Variations in environmental controls mean that the same data may be considered pseudonymised in one environment, but anonymous in another. See Box 5 and Section 3.2 on environmental controls.

3. “Anonymisation is impossible”.

Sometimes the term ‘anonymous’ is used by those outside of the world of data protection, for example in computer science, to mean data where the risk of re-identification is zero. This is, under some definitions, impossible if the data is also to be useful.

The reality? Crucially, this is not what the law requires. The law accepts some level of re-identification risk, meaning anonymisation is possible. It takes a risk based approach, rather than an absolute one, meaning there can be some potential for re-identification in anonymous data.

4. “Anonymisation = *k*-anonymisation or differential privacy, or...”.

A common refrain is that a particular approach will always result in anonymous data.

The reality? This is not the case. Both *k*-anonymity (Section 3.2) and differential privacy (Section 3.3) offer a way of thinking about privacy, they do not prescribe a given level of protection. They both have parameters, *k* and epsilon, that the organisation needs to set. The level of re-identification risk depends on that parameter (and other factors). Evaluating whether the data is in fact anonymous requires looking at what the re-identification risk is for that dataset with that parameter.


3. De-identification techniques



Identifiability is not a binary concept. It's more accurate to think of a sliding scale, where individuals are more or less identifiable depending on several factors. We'll use the term de-identification to mean reducing the likelihood that an individual is identifiable in a dataset. We'll refer to techniques that aim to identify an individual, or reveal information about them, as re-identification attacks.

De-identification techniques move data from left to right on the spectrum in our illustration below. Organisations can apply several de-identification techniques in combination to progressively reduce re-identification risk. The law aims to draw a line on that spectrum, considering data anonymous and out of the scope of the data protection regime once re-identification risk is low enough. As we explore in Section 4 below, the legal question of what level of re-identification risk is 'low enough' is challenging to answer.

This section explains the techniques commonly used to de-identify data and some re-identification attacks. It also explains the trade off between applying de-identification techniques and data utility and useability.

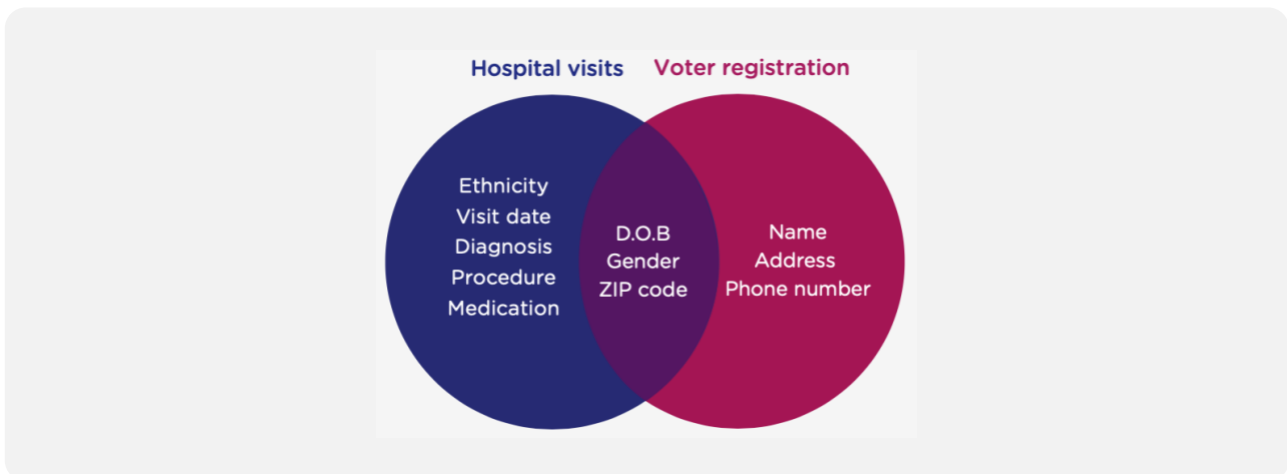


3.1. Masking, linkage attacks, and k-anonymity

Organisations often rely on masking direct identifiers to de-identify data. From a legal perspective this could result in pseudonymous data³, a subset of personal data under the GDPR (see Box 5). A typical approach involves two steps:

1. Selecting the values in a dataset that could, by themselves, single out an individual and allow data about them to be linked. These values are called “direct identifiers”. Some direct identifiers relate to just one person in all datasets (e.g. an email address or credit card number). Others may relate to just one person only in some datasets (e.g. a name may be unique in a small company’s employee data but not in population level data).
2. Once you select the direct identifiers, you can use a variety of techniques to mask the value. We’ve listed the most common masking techniques in Annex B, including deleting values, replacing values with tokens, and encrypting values.

Masking direct identifiers means that an individual is no longer directly identifiable, but Professor Latanya Sweeney’s pioneering [work](#) in the early 1990s showed that indirect identification is still possible. Sweeney focused on quasi identifiers, values that are not directly identifying on their own, but may be identifying in combination.



She examined data published by a hospital, containing the values in the blue circle in the Venn diagram. The data contained a patient’s date of birth, ZIP code and gender, which were often unique in combination and allowed an individual to be singled out. Sweeney demonstrated that if the same unique combination exists in another dataset, it can be used to link data relating to an individual.

³ The GDPR does not recognise the term pseudonymous data and instead refers to a process, that of pseudonymisation. We use the term ‘pseudonymous data’ in this paper for convenience, for example when contrasting it with the concept of ‘anonymous data’.

She obtained a voter registration dataset from a public records office and used the overlapping unique combination of values to link the two. This allowed her to associate a named individual with the hospital data. To prove that the data was identifiable and highlight the risk to individuals, she posted the then Governor of Massachusetts, William Weld, his medical records.

Sweeney went on to suggest a defence against this kind of 'linkage attack'. The linkage was only possible because 87% of the US population are likely to be unique on the basis of a five-digit ZIP code, gender and date of birth. In some cases, organisations can protect against linkage attacks by generalising data. For example, using just the year and month of birth instead of a full date of birth or a 3-digit ZIP code instead of a 5-digit ZIP code. By generalising data in this way you can reduce the precision in the dataset until no one is unique. With three-digit ZIP code, gender and month and year of birth the percentage of the US population likely to be unique falls to 0.04%.

Sweeney codified this in the idea of [k-anonymity](#). The k variable is the minimum number of people that are indistinguishable in a dataset. A dataset will be k -anonymous if each row is indistinguishable from $k-1$ other rows. $k=5$ means that each individual in the dataset would be indistinguishable from at least 4 others. Generalising the data more gives a higher number of k , and therefore a stronger privacy protection.

3.2. *k-anonymity, Big Data, and environmental controls*

When Sweeney developed k -anonymity in the late 1990s, many datasets looked like the hospital data in her linkage attack example; a few dozen fields relating to an individual. We're now in the data age and modern datasets, such as location traces or transaction histories, may contain thousands of data points on an individual. Applying k -anonymity to these types of data requires more generalisation, which can significantly reduce data utility.

To preserve data utility, organisations should consider combining controls on the data itself, like k -anonymity, with other types of controls. We've seen that linkage attacks require the use of auxiliary information (Sweeney used voter lists), the opportunity and motivation to carry out the linkage. Generally they do not happen by accident.

Environmental controls reduce the motivation and opportunity for an attacker to carry out a linkage attack, or to use the data in any other unauthorised way. Examples include requiring users to access the data in controlled settings (e.g. secure rooms, where their actions are monitored), carrying out background checks to vet data users before they can access data, requiring users to sign contracts with penalties for data misuse, or providing access to data in a way that does not allow it to be combined with other data (e.g. via a restricted portal, or a [trusted research environment](#)).

Environmental controls allow organisations to take a pragmatic and holistic approach to re-identification risk. Unlike some data transformations, environmental controls do not offer an easily quantifiable measure of residual risk. Generally speaking, environmental controls do not affect data utility - the analytical value of the data. However, they can make it harder to use data - reducing data useability - for example because of the expense or delay involved in applying them. For example, the cost of setting up a secure room or background checks delaying data access.

The history of linkage attacks, k -anonymity, and environmental controls illustrates four important concepts:

- De-identification comes at a cost. For example, generalising data makes it less identifiable, but can also reduce data utility. For example, generalising the date of a hospital visit to a year e.g. 4/7/2010 to 2010, will be acceptable for some purposes (e.g. calculating how many patients were seen in a year) but render the data useless for others (e.g. seeing which months are busiest). Likewise, environmental controls can affect usability. Secure locations can be expensive to run, and requiring data users to be vetted can delay projects. As with the impact on utility, how harmful this is will depend on the requirements of the project in question.
- Some de-identification techniques are effective for some datasets but not for others. For example generalisation can work well for demographic data, as Sweeney demonstrated. However, it can be less effective when the data is high dimensional, in other words when there are many data points about each individual, often over a period of time (e.g. transaction histories or location traces). k -anonymity is not a useful technique if there are too many quasi identifiers for it to provide both an acceptable level of privacy and useful data.
- The process Professor Sweeney went through: identifying an attack, demonstrating how it can be exploited, and developing a response highlights how the field has developed. We can apply the same approach, thinking about attacks and when they can be executed, to assess re-identification risk for a dataset.
- What can be used to re-identify someone is not intuitive, nor is what might be revealed. The state of Massachusetts didn't think that people could be re-identified in their de-identified health records. Nowadays that may seem more obvious, but the field has numerous examples of where those releasing data had made erroneous assumptions as the two examples below show.
 - In 2014 the New York City Taxi & Limousine Commission (TLC) released a dataset of de-identified taxi trips in response to a freedom of information request. Researchers examining the data highlighted two significant flaws in TLC's approach to de-identification. First, they used hashing to mask each taxi's unique medallion number. Hashing was not an appropriate technical method for this type of data and was "[trivial to undo](#)".

Another researcher [showed](#) that time-stamped photos of celebrities could be used as auxiliary information for linkage attacks because the taxi dataset included the time and location of a pick up or drop off. The photograph plus the timestamp revealed an individual's route, and other details about their trip, such as how much they tip.

- In 2007 Netflix released a large dataset of de-identified movie ratings. Researchers [used](#) publicly available ratings data from IMDB to re-identify the Netflix data. One affected woman [sued Netflix](#) claiming that her ratings history revealed her then hidden sexual orientation.

3.3. *Aggregates, reconstruction attacks and differential privacy*

The previous example considered row level data; datasets where each row relates to an individual. In contrast, aggregate data provides information about groups of individuals. We sometimes hear the terms 'anonymous' and 'aggregated' used interchangeably. However aggregates can be subject to re-identification attacks, meaning that not all aggregate data is anonymous. We'll consider one example of an attack on aggregates, and show how it relates to the legal definition of identifiability.

Reconstruction attacks rely on the fact that every statistic based on a dataset reveals some information about that dataset. In 2003 Professors Kobbi Nissim and Irit Dinur demonstrated that an attacker could work backwards from a set of statistics to [reconstruct](#) the underlying dataset. This is possible because for any set of statistics there is a limited (though potentially very large) set of possible underlying data that could have resulted in that statistic being true.

A simple example illustrates this: if I publish the fact that I have £2 worth of coins, an attacker can guess all of the possible combinations (around 74,000) of coins that make up that amount. As I release more information about the combination of coins I have (e.g. the number of 5p coins) the number of possible combinations is reduced, and an attacker's confidence that I have a specific coin (e.g. a 50p piece) increases.

This type of attack is problematic for organisations that publish large numbers of statistics based on personal data. The US Census Bureau carried out internal reconstruction attacks against the eight billion aggregate statistics derived from the 2010 Census. The aggregates were based on 25 data points on roughly 308 million individuals. The attackers were able to [reconstruct](#) data relating to voting age (yes/no), sex, age (in years), race and ethnicity (Hispanic or not). In 46% of cases (142 million people) the reconstructed data was an exact match for the underlying data.

In 2006 Nissim, in collaboration with Cynthia Dwork, Franck McSherry and Adam D Smith, developed a [method](#) to protect against reconstruction attacks by adding noise to statistics of the order of the contribution that any one individual would make. This can produce enough uncertainty to defend against reconstruction attacks, while still allowing for useful statistics. This approach is known as differential privacy, and has since been adopted by a number of organisations (including Google, Apple and the US Census). Differential privacy is not an algorithm itself, but a property of an algorithm. There are a variety of differentially private algorithms that can be used for different tasks, like producing aggregate statistics.

Despite its advantages, differential privacy will not be suitable where organisations need to use row level data. It can also be complex to deploy and requires expert analysis to ensure that an appropriate amount of noise is added to the data.

Deeper dive on differential privacy



Differential privacy gives a provable mathematical guarantee about privacy in terms of how much could be learned about a specific individual from the released data. That makes it unlike many other fields in privacy as it protects not against current known attacks but all attacks, those we know about and those we don't.

For more detail see Privitar's paper on differential privacy for the [UK's Office for National Statistics](#), and Privitar's appendix on anonymisation with differential privacy in the second edition of the [Anonymisation Decision-Making Framework](#).

The problem with 'significant' inferences



As explained in Section 6.2, one criterion used to assess whether data is personal or not is whether it allows for significant inferences to be drawn.

Consider two statements;

“John Smith has diabetes”

“7% of the UK population has diabetes”

Most people would be comfortable saying the first statement is personal data, and the second isn't. But what about statements in between?

Take the statement “25% of people living at 1 New Road Street have diabetes”. Perhaps John Smith is one of four people living in that house, so the statement reveals that one of those four people has diabetes. The statement doesn't reveal John Smith's health status definitively, but it increases an attacker's confidence that John Smith has diabetes. You can imagine a range of other scenarios that would alter someone's confidence that John Smith has diabetes by differing amounts. Which of these is personal data and which isn't?

In today's data decision driven world, data doesn't need to be definitive for it to be used to make meaningful decisions about people. In fact, due to the messy nature of data, decisions often can only be made on likelihood, as opposed to certainty. As such it would seem wrong to say only definitive statements constitute personal data, but then where is the line?

Some have suggested concepts such as plausible deniability or when the inference is significant, but these don't tell us practically how to make an evaluation. The problem is analogous to the heap paradox, whereby grains of sand are taken from the heap until there are none left. At which point did it cease to become a heap? In the same way, when does an inference that can be drawn from a statement cease to be significant?

3.4. Machine learning models and model inversion attacks

With the rapid increase in use of machine learning in recent years, some organisations share models, rather than the training data that the models were trained on. At first glance one might assume that the models themselves aren't disclosive. However, like aggregated data, they can in some instances [leak information](#) about specific individuals. What this means for the legal status of the data, specifically whether models themselves can be classified as personal data, is debated.

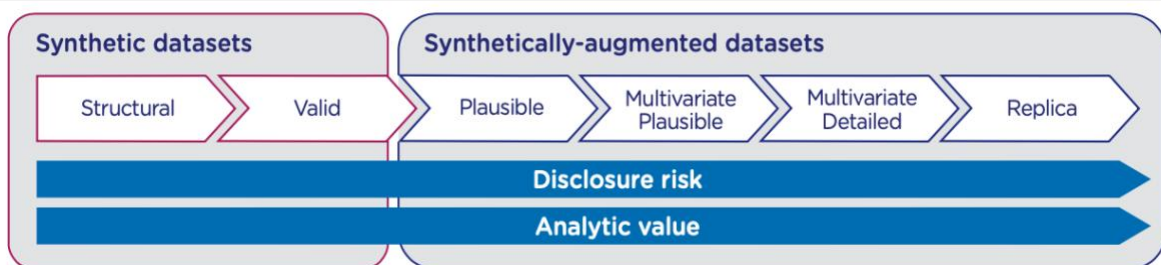


An image recovered using a new model inversion attack (left) and a training set image of the victim (right). The attacker is given only the person's name and access to a facial recognition system that returns a class confidence score.

While research in this area is still emerging, fundamentally models record some of the personal data that they are trained on, and in some situations can be made to reveal that data. One example is called a model inversion attack. An attacker with some information about someone in the training data, can get the model to reveal additional information. For example, by looking at the confidence scores in a facial recognition model, researchers were able to [reconstruct](#) the face of an individual in the training dataset shown below. The same approach can be used to reveal other types of data.

3.5. Synthetic data

Another approach to providing useful and privacy preserving row level data, is to create synthetic data. Synthetic data is an umbrella term, covering a range of approaches for generating data with varying levels of utility. The ONS [proposed](#) a scale based on resemblance to raw data, purpose and disclosure risk. As you can see, the higher the utility in the synthetic data, the greater the risk.



There has been significant recent interest in a new approach to creating synthetic data using Generative Adversarial Networks (GANs). GANs work by having two models compete; using a real dataset as the input, one model generates a synthetic dataset, and the other tries to assess if it is real or synthetic, in this way the two models compete, leading to more realistic looking synthetic datasets over time. This approach is used to generate realistic 'photographs' showcased on thispersondoesnotexist.com.

Research on GANs, and other types of synthetic data, is nascent. It is unclear what types of vulnerabilities exist in GAN-generated synthetic data. For example, could the GAN-generated synthetic data itself be subject to something similar to a model inversion attack? The lack of certainty around vulnerabilities makes it difficult to assess re-identification risk associated with GAN-generated synthetic data. This creates uncertainty as to whether and when that GAN-generated synthetic data could be considered personal data.

There are also significant questions around the utility of GAN-generated synthetic data. For example, data analytics use cases aim to identify trends in data to inform business decisions. In some cases, it may be difficult to separate a trend which actually exists in the underlying data from one which was created artificially by the GAN. Organisations considering synthetic data may want to [contact us](#) for more detail.

3.6. Balancing competing priorities

The discussion of technical and environmental controls in this section highlights the need for organisations to balance competing priorities: privacy or re-identification risk, data utility and data useability. Controls reduce re-identification risk but can also affect data utility and useability. Using controls effectively requires organisations to 'fine tune' each control to deliver data suitable for the intended purpose.

This links back to our introductory comment that anonymisation may be difficult, or impossible, in the sense that anonymous data may not deliver the utility and usability required for a specific use case. Section 4 below on the legal position describes the risk assessment an organisation should carry out to determine whether data is anonymous.

We also highlight uncertainty in the legal position with regard to environmental controls, specifically how and whether they can be taken into account when assessing re-identification risk (Section 5). Excluding environmental controls from consideration makes it more difficult to achieve anonymous, high utility data.

From a technical perspective we can conclude that:



Organisations can manage or mitigate re-identification risk through data transformations, environmental controls or (more likely) both in combination.



Applying controls, whether by transforming the data or controlling the environment, has an impact on data utility and usability. In some cases, the controls necessary to anonymise data will affect utility and useability to the extent that the data is no longer suitable for the intended purpose.




Identifiability and re-identification risk are not intuitive. Real world examples show it is easy to get wrong. See the examples in Section 3.2.



By thinking about what attacks could be effective, we can reason about the likelihood of re-identification.

4. Anonymisation: the legal position in the UK

Our analysis of anonymisation from a legal and public policy perspective draws on legislation, guidance from regulators, and UK and European case law.



4.1. Legislation

The legal definition of anonymisation comes from Recital 26 of the GDPR:

“The principles of data protection should apply to any information concerning an identified or identifiable natural person. Personal data which have undergone pseudonymisation, which could be attributed to a natural person by the use of additional information should be considered to be information on an identifiable natural person.

To determine whether a natural person is identifiable, account should be taken of all the means reasonably likely to be used, such as singling out, either by the controller or by another person to identify the natural person directly or indirectly. To ascertain whether means are reasonably likely to be used to identify the natural person, account should be taken of all objective factors, such as the costs of and the amount of time required for identification, taking into consideration the available technology at the time of the processing and technological developments.

The principles of data protection should therefore not apply to anonymous information, namely information which does not relate to an identified or identifiable natural person or to personal data rendered anonymous in such a manner that the data subject is not or no longer identifiable. The Regulation does not therefore concern the processing of such anonymous information, including for statistical or research purposes.” (underlining and paragraph breaks added)

Taken together, the first and third paragraphs distinguish between personal data, which is information concerning an identified or identifiable natural person, and anonymous information, which does not. They state that anonymous information is out of scope of the data protection regime.

The second paragraph describes the boundary between personal data and anonymous information. It is based on whether an individual is “identifiable” taking account of the “means reasonably likely to be used” to identify an individual. As we discuss below, the ICO’s ‘motivated intruder’ test provides a methodology for assessing whether an individual is identifiable.

The third paragraph allows for the scenario where information that is personal data is “rendered” anonymous. The text does not provide any guidance on how data can be rendered anonymous but clearly envisages that it should be possible. The formulation in the third paragraph is notably circular: data protection does not apply to “...personal data rendered anonymous in such a manner that the data subject is not or no longer identifiable (sic)”.

Pseudonymisation

The GDPR defines pseudonymisation in Article 4(5). Pseudonymisation means “the processing of personal data in such a manner that the personal data can no longer be attributed to a specific data subject without the use of additional information, provided that such additional information is kept separately and is subject to technical and organisational measures to ensure that the personal data are not attributed to an identified or identifiable natural person”.

Pseudonymisation requires three elements. Removing direct identifiers from the data, keeping those identifiers separately from the data, and subjecting them to technical and organisational controls. In practice this could mean:

- Tokenising or encrypting a direct identifier, essentially removing it from the data
- Keeping the token mappings or encryption key separately from the data
- Protecting the token value or encryption key with
 - Technical measures (e.g. cyber security controls)
 - Organisational measures (e.g. policies or standard operating procedures banning the unauthorised use of the token vault or encryption key).

Data that has been pseudonymised is still personal data under the GDPR. Pseudonymisation is a safeguard, as the GDPR makes clear; where considering whether processing for a purpose other than that for which the personal data have been collected is compatible with the original purpose, account should be taken, amongst other things, of “the existence of appropriate safeguards, which may include encryption or pseudonymisation⁴.” It reduces the risk to individuals but, unlike anonymisation, does not take data out of the scope of data protection law.

⁴ Article 6(4)(e) GDPR


Pseudonymisation (cont.)

This gives us three general categories; directly identifying data, pseudonymous data, and anonymous data. Crucially, both of the first two are personal data. It's important to note that the boundaries between these categories can be blurred, and within each data can be more or less identifying.

For example, a dataset containing names is directly identifying, but not as identifying as a dataset containing names and contact details. Likewise, a dataset with name and contact details removed, but postcode and date of birth retained is more identifying than the same dataset with dates of birth removed, even though both could be pseudonymous.

As shown below, identifiability is a spectrum. The legal status of data changes as you move along the spectrum. As you can see, the spectrum shows a grey area between pseudonymous and anonymous, indicating uncertainty as to where this boundary is.

De-identification techniques make re-identification less likely



Re-identification attacks make data more identifying

There can be big differences within each section of the spectrum. For example, data that has had all direct identifiers removed, but still has a lot of quasi identifiers (such as date of birth, or postcode), would be at the leftmost section of the pseudonymous category. This would present a very different risk to data that is at the rightmost edge, near the border with 'Anonymous'. As you move to the right of the spectrum, even within the 'Pseudonymous' section, the potential number of lawful uses for the data increases.

4.2. Guidance from regulators

The available general guidance on anonymisation in the UK comes from the Information Commissioner's Office (ICO) [anonymisation code of practice](#), published in 2012. The ICO also funded the establishment of the [UK Anonymisation Network](#) (UKAN), to advance anonymisation best practices and offer more detailed practical advice on anonymisation. UKAN published the second version of the [Anonymisation Decision Making Framework](#) (ADF) in 2020⁵.

At the European level, two Article 29 Working Party (A29WP)⁶ opinions are of particular relevance: [Opinion 04/2007](#) on the concept of personal data and [Opinion 05/2014](#) on anonymisation. As we'll explore in Section 6 below, the ICO and the A29WP took different approaches on some key issues.

The ICO's 2012 guidance describes the methodology for assessing re-identification risk based on the 'motivated intruder' test, which we describe in Section 4.4 below. The A29WP guidance focuses on anonymisation techniques. Both sets of guidance pre-date the GDPR. The ICO guidance is based on the UK Data Protection Act 1998 and the A29WP guidance on the EU's 1995 Data Protection Directive.

Both the ICO and A29WP guidance are pending updates. The ICO [recognises](#) that "questions about when data is personal data or anonymous information are some of the most challenging issues organisations face" and has outlined plans for refreshed guidance. The EDBP's [work plan](#) for 2021/22 also promises guidelines on anonymisation and pseudonymisation.

Sector-specific approaches to anonymisation



Anonymisation is important for, and widely used in, health research and clinical data sharing more generally. Guidance on anonymisation in scientific research covers familiar ground regarding identifiability and assessment of likely intruders but it also introduces scenarios unique to the sector, where anonymising data will not be possible, such as where NHS numbers are to be left in the data.

Another unique feature of the sector in the UK is the common law duty of confidentiality in relation to medical personal data. This is partially outside the jurisdiction of data protection authorities when assessing the lawfulness of processing of medical data and questions of identifiability, leading to confusion and a current lack of clear, overarching regulatory guidance about the interaction between the application of common law rules and GDPR rules as they relate to anonymising such data.

⁵ Privitar helped in the drafting of the ADF by authoring the appendix on anonymisation and differential privacy.

⁶ The predecessor to the European Data Protection Board (EDBP)

Sector-specific approaches to anonymization (cont.)

Regulatory bodies and health research standards setters have sought to address these issues, and issue their own, sector-specific, guidance on anonymisation.

For example:



NHS Digital's [Anonymisation standard](#) for publishing health and social care data, which is also pending an update.



The Medical Research Council's [Guidance note 5](#) on identifiability, anonymisation and pseudonymisation. The guidance highlights unique features of datasets such as genetic data, where identification is likely to have an increased impact on family members and not just the individual themselves. Bristows' summary of the guidance is [here](#).



The European Medicines Agency (EMA) [Policy 0070](#) (2019), associated [guidance](#) on implementation and the EMA's [workshop report](#) on 'data anonymisation - a key enabler for clinical data sharing' (2018).

4.3. Case Law

To complete our analysis, we examined relevant UK and European case law. The key cases, with commentary, are listed in Annex A below. The UK case law needs to be treated with caution when considering anonymisation in a practical context because:

- Most of the case law is based on the UK's Data Protection Act 1998, which implemented the 1995 EC Data Protection Directive. There are differences between the definitions of personal data in the two, and furthermore a widening in the GDPR of the definition in the 1995 Directive. See Section 5.1 for more detail on the differences between the definitions.
- Much of the case law relates to freedom of information requests. Under UK freedom of information laws, no conditions (whether contractual, technical, or environmental) may be imposed on recipients in order to control the information once released. In contrast, many organisations are looking for options to use data internally, or share it in a controlled way, rather than making it, in effect, publicly available.

4.4. When is data anonymous?

The current UK position is based on the ICO's 2012 guidance, which:

- Describes an attack-based risk assessment called the 'motivated intruder' test that includes a set of assumptions about the attacker's capabilities and motivation.
- Suggests that a motivated intruder should be considered successful if they would be able to identify an individual.
- Accepts some level of re-identification risk. The motivated intruder must have a reasonable likelihood of success.

Distinguishing 'identifiable' and 'named'



The ICO guidance is clear that an individual can be identifiable without being named. In other words, an individual is identifiable if an attacker can establish a reliable connection between an individual and some data relating to them even if the attacker is not able to name the individual in question.

Assessing identifiability means assessing the probability that an individual could be identified. The law sets a "reasonable likelihood" test. If it is reasonably likely that an individual could be identified from the data, that data is considered personal data. This section summarises our analysis of the current case law and guidance on carrying out the 'motivated intruder' test, also called the 'determined person' test in recent Tribunal decisions. The ICO's 2012 guidance states that the motivated intruder test "involves considering whether an 'intruder' would be able to achieve re-identification *if* motivated to attempt this" (emphasis in the original text).

We can think of the 'motivated intruder' described in the ICO's 2012 guidance as one attacker profile. The case law describes other attacker profiles, with different assumptions about their capabilities. Taken together, the case law and guidance suggest that organisations should always consider the motivated intruder attacker profile as a baseline, and make a case by case assessment of whether other, more capable, attacker profiles are relevant.

Applying the test requires organisations to consider:

- **Which attackers would be likely to attempt re-identification?** The guidance and case law describe several attacker profiles and suggest that some data will be more attractive to particular attackers. The ICO guidance notes that some data may be “innocuous” on the face of it, but that organisations should still carry out “a thorough assessment of the threat of re-identification”. However, as we explore in Section 5.2, some case law has questioned whether an assumption that all de-identified data will be subject to attack leads to an overly risk averse approach.
- **What capabilities would each attacker possess?** The test allows for a set of assumptions about the attacker’s capabilities. This relates to the “means reasonably likely to be used” element of the GDPR Recital 26 definition. Taken together, the case law and guidance define a set of capabilities organisations should consider:
 - The ICO’s 2012 guidance states that organisations should assume that the motivated intruder is “reasonably competent”. This is a higher bar than asking whether an “inexpert” member of the public could achieve re-identification but a lower bar than considering whether someone with “access to a great deal of specialist expertise, analytical power or prior knowledge” could do so.
 - Summarising the test, the court in *Peters* explained that the motivated intruder will take reasonable steps to achieve re-identification, be determined and competent and have access to resources such as the internet, libraries and public documents.
 - The courts have considered other attacker profiles. For example, an investigative journalist, with the journalist assumed to have additional skills (e.g. in *Magherafelt DC* and in *Craigdale*). *Department of Health* considered other groups, such as campaigners.
 - Guidance and case law are clear that the motivated intruder will not resort to criminality (e.g. hacking or burglary to obtain protected information). The court in *Peters* noted that “a motivated intruder is not a lawbreaker”.

Privatar’s recommendations to the ICO on anonymisation include examples of five attacker profiles. Organisations may need to define the attacker profiles relevant to the data in question and assess the data against each of those profiles. In some situations it may be necessary to consider stronger attackers than case law has covered.

The test will often need to be applied iteratively. The organisation defines the set of attackers and their properties, applies controls and checks what effect the controls have on the likelihood of a successful re-identification attack, repeating the two latter steps until an acceptable balance between utility and re-identification risk is achieved.

The ICO guidance is also clear that the assessment is not a one off exercise. Changes in the context, such as new auxiliary information becoming available or new types of re-identification attack being discovered, may change the likelihood of an attacker carrying out a successful re-identification attack.

The relationship between controls and the motivated intruder test

Organisations can reduce the likelihood of a successful re-identification attack by applying controls to the data and to the environment in which the data is processed, de-identification techniques and environmental controls respectively. We discuss de-identification techniques in Section 4 above, with more detail at Annex B. We also list common environmental controls at Annex B.

For example, controls can:



Reduce the motivation to attack data. Controls on data could include record flagging, to ensure that records relating to high profile individuals are removed from the de-identified data. This could reduce the motivation for an investigative journalist to attack the data. Environmental controls, such as access logging, can also reduce motivation by increasing the likelihood that an attacker's actions would be recorded, detected, and acted upon.



Limit the opportunity for an attack. Controls on who can access the data and under what conditions would have a significant effect on limiting the set of possible attackers. For example, requiring de-identified data to be used on site in a specific secure room limits access to people with the necessary credentials.




Make the data less 'identifying'. For example, restricting an attacker's ability to link the de-identified data with other information by applying technical controls (e.g. k -anonymity) to protect against linkage attacks.

This way of thinking about attackers and the legal threshold can guide decisions about which controls to apply and how to configure those controls (e.g. if using k -anonymity, what value of k is appropriate?).

5. Summary of issues with the UK approach

5.1. The definition of personal data

This section highlights four of the key issues or challenges with the UK approach to anonymisation. They arise from inconsistencies within the case law, or between case law and guidance.



Comparing definitions of personal data



The definition of personal data in the DPA 1998 is in section 1:

“personal data” means data which relate to a living individual who can be identified —

(a) from those data, or

(b) from those data and other information which is in the possession of, or is likely to come into the possession of, the data controller, and includes any expression of opinion about the individual and any indication of the intentions of the data controller or any other person in respect of the individual;

The definition in Directive 95/46/EC is in Article 2:

(a) 'personal data' shall mean any information relating to an identified or identifiable natural person ('data subject'); an identifiable person is one who can be identified, directly or indirectly, in particular by reference to an identification number or to one or more factors specific to his physical, physiological, mental, economic, cultural or social identity;

For completeness, the definition in the GDPR, article 4(1) is:

'personal data' means any information relating to an identified or identifiable natural person ('data subject'); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person;

We set out in Section 4 the GDPR approach to anonymous information. As also noted above, the UK changed the definition of personal data in its Data Protection Act 1998 from that in the EC Data Protection Directive, just as some other EU Member States did in their national implementing legislation.

The definition of personal data in the 1998 Act introduced a condition which was not present in the 1995 Directive and is not in the GDPR. The 1998 Act defined personal data as “data which relate to a living individual who can be identified (a) from those data or (b) from those data and other information which is in the possession of, or is likely to come into the possession of, the data controller”.

Part (b) of the definition in the 1998 Act places a limit on the other information that could be used to re-identify an individual, namely that the “other information” should be in the possession of or likely to come into the possession of the organisation.

The GDPR Recital 26 adopts a broader approach, referring, in relation to identifiability, to “all of the means reasonably likely to be used...by the controller or any other person [to identify a natural person]”. This includes the possibility of relying on information that is not in, and is not likely to come into, the possession of the controller. Controllers undertaking a re-identification risk assessment must take this key difference between the definitions into account in their assessments.

The court in *Peters* highlights the difference in definitions of personal data. The case was decided on the basis of the 1998 Act, because the GDPR was not yet in force when Mr Peters brought his claim. However, the court explicitly considered whether the decision on whether the data in question was personal data would have been different under the GDPR definition and concluded that it would not have been. This was not, however, a foregone conclusion in other cases.

Several of the cases under the 1998 Act appear to struggle with the question of whether re-identification of a natural person *by the data controller* is relevant to the assessment of identifiability in the case at hand. The key cases do not reach the same conclusion about the approach.

Identifiable from whose perspective?

It is important to recognise this apparent difference of approach in the case law to the assessment of identifiability. The *Common Services Agency* case provides (per Lord Hope's leading judgment) that the assessment is of whether data are identifiable by the controller as well as third parties.

The later *Department of Health* case, on the other hand, provides that data may be personal in the hands of the controller but de-identified in the hands of recipients. This uncertainty gives controllers an additional hurdle in assessing identifiability, and with it re-identification risk. See Sections 1 and 3 of Annex A for more detail.

The European case law on the issue considers identifiability in the round, rather from the perspective of any particular party. The approach in *Breyer* notably also calls into question the suggestion in the A29WP's Opinion 05/2014 that data *cannot* be anonymous in the hands of the recipient where the discloser of the data retains the original personal data and may accordingly be able to re-identify the anonymised data disclosed to the recipient. This is further discussed in Section 5.3.

In *Nowak*, the CJEU was asked to consider whether "the written answers submitted by a candidate at a professional examination and any examiner's comments with respect to those answers constitute personal data".

The CJEU held that it was irrelevant whether the examiner could identify a candidate. It was sufficient that anyone, in this case the exam board, could identify the candidate. The CJEU, referring to *Breyer*, concluded that "there is no requirement that all the information enabling the identification of the data subject must be in the hands of one person".

The CJEU's approach in *Nowak* broadly aligns with Lord Hope's judgement in *Common Services Agency* and holds that organisations must consider identifiability from multiple perspectives, not solely from the perspective of one party, such as the data recipient.

5.2. Is there a realistic possibility of an attacker?

The case law is unclear as to whether organisations should consider whether there is a realistic possibility of an attacker separately to the broader question of re-identification risk. *Peters* and *Miller* separate out the two issues, in contrast to *Department of Health*, where Mr Justice Cranston combines them.

Peters concludes that "there has to be an assessment of whether there is at least a realistic possibility of a motivated intruder in a particular case". However, the Tribunal did not actually carry out that assessment and simply proceeded on the basis that there was a realistic possibility of a motivated intruder. The Tribunal noted that carrying out this assessment would reduce the risk that a requester "might be denied information to which he or she would otherwise be entitled simply because of a risk which was not grounded in reality".

The judge in *Miller* took a slightly different approach. He found that the nature (statistics on homelessness) and age (several years old) of the data in question meant that it was unlikely that anyone would be sufficiently motivated to attempt to re-identify the individuals represented in the de-identified data.

Despite concluding that the data was unlikely to be of interest to an attacker, the judge went on to consider the risk of re-identification. *Miller* suggests that, in line with the ICO guidance, simply concluding that the de-identified data would not be of interest to an attacker does not excuse the organisation from assessing re-identification risk.

In *Department of Health*, Mr Justice Cranston held that assessing the likelihood of re-identification included considering “...the likelihood that particular groups such as campaigners and the press will seek out information of identity...” This approach treats the issue of whether an attack is a realistic possibility as just one element of the broader question: whether an attacker could successfully re-identify an individual in the data.

Our view, in light of the inconsistent approach in the case law, is that the right approach is to assess re-identification risk in the round, that is, per *Miller*, even if the conclusion is formed that the de-identified data would not be of interest to an attacker, the rest of the assessment (of re-identification risk from a motivated intruder) remains essential.

In doing so, the organisation must exercise *flexibility* in its risk assessment. Set the bar too high, and the organisation may impose unnecessary constraints on its own freedom of use of the data in question, or that of others seeking access to it. Set it too low, and the data will not be lawfully anonymised.

The case law also shows that the courts have grappled with the issue of overzealous risk avoidance as a barrier to access to information, under freedom of information laws. An overly cautious approach constrains access to data, limiting research and other matters in the public interest, including genomics and health-related research and the publication of useful statistical and other information that, under freedom of information laws, the public has a right to know.

Lawmakers are aware of these challenges. For example, see the Department of Health and Social Care [White Paper](#) on legislative proposals to build on the NHS Long Term Plan, analysed in this Bristows [paper](#). Similarly the [Caldicott Principles](#) note that “good information sharing is essential for providing safe and effective care” and that “there are important uses of information for purposes other than individual care, which contribute to the overall delivery of health and social care or serve wider public interests”.

Attempts to balance the public interest in access to information (the requesters in both *Peters* and *Miller* were researchers) with the requirements in data protection law has created inconsistencies in the case law. In addition, most case law involves releasing information publicly in response to freedom of information requests. This means that the courts have had less opportunity to deal with questions on the role that environmental controls can play in reducing re-identification risk.

How are you managing the trade offs?



The ‘public good’, (for example, scientific research), however, is not a relevant legal factor in assessing re-identification risk. In our Conclusion, we invite examples from readers of situations for further investigation where a tension arises between, on the one hand, the protection of personal data, and on the other, another public good (whether in relation to scientific research, the public’s ‘right to know’ under freedom of information laws or another field (see Section 7)). Bristows’ introduction to some of the issues in the life sciences sector can be found [here](#).

5.3. *How to account for the data environment when assessing re-identification risk?*

As discussed above, most of the case law considering the ‘motivated intruder’ test relates to publishing information under UK freedom of information laws. Under these laws, no conditions (whether contractual, technical, or environmental) may be imposed on recipients in order to control the information once released.

The information to be released must be assumed to enter the public domain without restriction. As such, these cases can be distinguished from the scenario where data is shared in a controlled manner, that is, with contractual, technical and/or environmental conditions attached. The ICO’s 2012 guidance notes that re-identification risk “will vary according to the local data environment and particularly who has access to information”.

This suggests that environmental controls are relevant to assessing re-identification risk (certainly, at any rate, outside the freedom of information context), but does not provide any further guidance on what types of controls are relevant. UKAN [recommends](#) that organisations consider the ‘data situation’. The ‘data situation’ approach considers the “aggregate set of relationships between some data and the set of their environments.”

European case law touches on but does not provide conclusive guidance on environmental controls, insofar as the legal framework can be said to constitute an ‘environmental’ control or factor. In *Breyer* the European Court of Justice found that dynamic IP addresses held by the German government were personal data because the government “has the legal means which enable it to identify the data subject with additional data which the internet service provider has about that person”. The Court noted that the German government could, in some circumstances (e.g. a cyber attack), use legal powers to obtain the necessary information from the internet service provider.

Breyer's perhaps most helpful contribution to the issues discussed in this paper is the recognition of identifiability as a relative and not an absolute concept. *Breyer* held that it was not reasonable to assume that the online service provider in the case would approach *any third party* holding additional information but instead that the Directive must have been referring only to those third parties whom the service provider might approach *in a reasonable manner*. In doing so, the case provides a useful counterweight to the suggestion in the A29WP's Opinion 05/2014 that data cannot be anonymous in the hands of the recipient where the discloser of the data retains the original personal data and may accordingly be able to re-identify the anonymised data disclosed to the recipient.

Breyer shows that the overall context is relevant to assessing re-identification risk. This raises the question of what other factors, beyond the data itself, may be relevant to that assessment. In contrast, the A29WP's Opinion 05/2014, considers only the data itself and not other factors, such as the overall context of the processing activity (or, in the words of UKAN, the 'data situation') as relevant to the assessment.

Similarly in *Vidal-Hall*, data held within an organisation (in this case browser generated information (BGI) held by Google) was deemed to be identifiable because the same organisation held other information needed to identify that data. In *Vidal-Hall* the fact that the BGI and the other information were held separately and Google argued that it had not in fact, nor did it intend to, combine the two sets of information were not seen as sufficient safeguards in the circumstances, and thus ultimately not determinative in the re-identification risk assessment. Although it is doubtful that an internal policy banning data linkage would have changed the outcome, the lack of certainty in the application of this reasoning to more complex types of corporate structure (e.g. public health services) leaves uncertain the impact of this type of environmental control on re-identification risk.

5.4. According to the guidance, when is a re-identification attack considered successful?

This section describes the guidance on identifiability in the ICO's 2012 code, to the extent that it describes a successful re-identification attack. The code is not legally binding, so should be read alongside the case law discussed in Annex A.

The ICO's 2012 code suggests that a successful attacker must learn something 'new' about an individual. Framing success in this way is helpful for a risk-based approach to assessment of identifiability. The example the guidance provides is an individual's genetic code stating: "This [genetic code] would identify an individual uniquely, but only to someone who already has access to both the code and the identity of the individual it belongs to."


The example in our view has not stood the test of time. The GDPR of course recognised genetic information expressly for the first time as falling within the definition of personal data. A whole human genome sequence and indeed subsets would potentially be highly identifying, meaning something new could be learnt should someone with only an extract (enabling them to re-identify an individual) obtain the full code (enabling them to learn something new).

Unlike most types of data, this type of data is an example of one where utility generally increases over time (particularly bearing in mind the typical premise of research projects in this field, to gather growing amounts of data over time). That said, we welcome the acknowledgement that ‘success’ in relation to re-identification under the motivated intruder test in the ICO’s view involves the requirement to learn something ‘new’.

A key distinction between the UK guidance and that of the A29WP is that the latter tries to provide a (problematic) test for identifiability, based on singling out, linkability, and inference. The ICO Code doesn’t include this test, but instead offers a series of factors or criteria to consider. As to the UK case law, the reasoning is often so reliant on the context that it can be difficult to apply the principles more generally. Consequently, it does not always give very much clarity to organisations which find themselves in different scenarios. There are inherent contextual aspects of identifiability that may prevent a simple test from working. However, as this paper indicates, we consider that more effective signposts are achievable to make practitioners’ lives easier as to the types of factors that should be taken into account in making this assessment.

6. Divergence between the UK and EU positions

This section focuses on three key areas of divergence between the UK and EU legal positions. We refer to the latter as the EU position, recognising that this may obscure nuances in the positions of individual EU Member States.



The issues discussed in Section 5 and Annex A highlight gaps or inconsistencies within UK case law and between case law and guidance. The UK position also diverges from the EU approach in several areas.

Divergence has a significant impact on organisations operating across multiple jurisdictions, as different regulators may reach different conclusions as to whether data is anonymous.

The A29WP opinion on anonymisation recognises that national positions diverge. It states that:

Anonymity is interpreted differently across the EU – corresponding in some countries to computational anonymity (i.e., it should be computationally difficult, even for the controller in collaboration with any party, to identify directly or indirectly one of the data subjects) and in other countries to perfect anonymity (i.e. it should be impossible, even for the controller in collaboration with any party, to identify directly or indirectly one of the data subjects). Nevertheless, “anonymisation” corresponds in both cases to the process by which data are made anonymous. The difference lies in what is considered as an acceptable level for the risk of re-identification.

This recognition of the relevance of risk in assessing anonymity is helpful. Legally speaking, there is no consensus in the EU that anonymisation is an absolute state, in some way ‘scientifically’ or computationally ascertainable. Our view is that such an approach would be legally incorrect in the UK, and arguably also in the EU, based on the face of the EU GDPR. Rather, a risk-based approach is inherent in the assessment. A risk-based approach is compatible with the approach underpinning the GDPR generally.

6.1. Requirement to delete the source data in order to create an anonymous extract

Can a data extract be anonymous if the disclosing organisation retains the ability to identify an individual in the data? For example, if an organisation shares an extract of a dataset with a third party, can that extract be anonymous to the third party even though the originating organisation retains the ability to re-identify the extract?

UK guidance allows organisations to assess identifiability from the perspective of the data recipient. The ICO's 2012 guidance states that "where an organisation converts personal data into an anonymised form and discloses it, this will not amount to a disclosure of personal data. This is the case even though the organisation disclosing the data still holds the other data that would allow re-identification to take place".

The guidance is in line with some of the UK case law, for example in *Department of Health*, where the court considered the public policy impact and concluded that a requirement to delete the source data could undermine an organisation's ability to publish useful information (in this case statistics relating to health outcomes). See Annex A for a discussion of the ways in which the UK courts have grappled with this notion.

In contrast, the A29WP opinion states that a de-identified data extract remains personal data unless the controller deletes the original, identifiable data. Retaining the original, identifiable data means that the organisation retains the ability to re-identify an individual in the de-identified data extract.

This is crucially important in a business context. Many use cases involve producing a data extract from a 'live' or 'production' system. A requirement to delete the data from the live system in order for the extract to be considered anonymous could be unworkable in practice. For example, consider health data extracted from an electronic patient record. If anonymisation required deleting the source data this would clearly be unworkable in practice as the record will be necessary for the patient's care.

6.2. Identifiability: the regulators' perspective

The A29WP opinion sets out a three part test for identifiability. Data is identifiable if an individual can be singled out, data relating to them linked, or information about them can be inferred. We consider singling out and linkage in detail above (see Section 4.2). The A29WP opinion defines an inference as "the possibility to decide, with significant probability, the value of an attribute from the values of a set of other attributes".

The specific challenges related to inferences are addressed in Privitar's [anonymisation recommendations](#) to the ICO. Restrictions on inferences are arguably incompatible with deriving useful information from an anonymous dataset, as any useful information (i.e. information that allows the recipient to learn something new, even about a group) could increase the recipient's confidence about an inference.

The ICO's 2012 guidance, unlike the GDPR, does not use the term 'singling out'. It refers to linkage on a number of occasions, and focuses on the motivated intruder's ability to leverage other sources of information to achieve re-identification. The guidance makes a passing reference to inferences, in the context of a recommendation on controlling inferences via statistical disclosure control techniques (e.g. aggregation or small count suppression).

The ICO's 2012 guidance includes the concept of an intruder being able to learn something new about an individual from de-identified data. However, this is not explicitly discussed in the case law.

6.3. The threshold for identifiability; the regulators' perspective

The A29WP opinion states that "anonymisation results from processing personal data in order to irreversibly prevent identification". This indicates an absolute threshold, with no tolerance for re-identification risk, unlike the UK case law, the ICO's guidance and the GDPR.

In contrast, the ICO's 2012 guidance takes a risk based approach and assesses re-identification risk based on the 'motivated intruder' test, while UK case law confirms that re-identification risk does not have to be completely eliminated for data to be considered anonymous. For example, in *Department of Health* the court agreed with the Tribunal's finding that the data in question was not personal data because the risk of re-identification was 'remote'. There is support in the case law for going further and suggesting that to set re-identification risk at zero is to misinterpret the legal position, which is framed by reference to a test of likelihood.

There will of course be scenarios where it is the prerogative of organisations controlling a given dataset to set the bar higher than is legally required. Equally, we can expect many continuing arguments in the future where it is contended that a risk threshold involving elimination of the risk of re-identification is an unnecessary and counterproductive brake on, for example, research and development activities that may actually threaten the public good in the long term, by impeding or altogether stopping such research.


Key Questions arising from the guidance

The following table inevitably sacrifices some of the nuances discussed above in order to capture the key questions arising from our analysis of the guidance, and should be read in conjunction with that analysis.

| | ICO's 2012 guidance | Open questions relating to the ICO guidance? | A29WP opinion |
|--|--|---|--|
| Definition of personal data | Based on the UK's 1998 Act | The definition in the 1998 Act is different to that in the 1995 Directive and in the GDPR. | Based on the 1995 EC Directive |
| Which attacker profiles are likely to attempt re-identification? | ICO focuses on the motivated intruder. Case law gives further examples of attackers: the general public, investigative journalists, ex-partners, spies | How should an organisation determine which attacker profiles are relevant? What assumptions about capabilities or controls to limit them are relevant? Case law suggests a separate assessment of whether an intruder is likely to be motivated to attack the data | Re-identification risk is considered for any and all potential attackers |
| How to account for the data environment? | Controls on the environment are relevant to assessing re-identification risk | Lack of clarity arising from <i>Vidal-Hall</i> | No discussion of environmental controls |
| What does success look like for the motivated intruder? | Focus on linkage | To what extent does the ICO's 'learning something new' test overlap with the A29WP's 'inference'? | Singling out, linkage or inference |

7. Conclusion - What does this mean for organisations?

Organisations seeking to use data for secondary purposes have several options. There is no one-size-fits-all approach, as the most appropriate option depends on the specific context. Organisations have to balance re-identification risk, utility and useability to ensure that data remains suitable for the intended purpose.



As this paper shows, anonymisation is one option. Anonymisation takes data out of the scope of data protection law, meaning that it can be used without assuming the obligations arising from data protection law. Note, however, that the act of anonymising the data is a processing operation that falls within the scope of the GDPR, which must be taken into account when considering this option. Furthermore, anonymisation is not a silver bullet. It can be challenging to achieve and could carry a significant cost in terms of data utility and usability.

The GDPR offers other options. In particular, as discussed in call out Box 5, de-identification can increase the breadth of lawful use cases by reducing the risk associated with the processing. The further towards 'anonymous' on the identifiability spectrum it is, generally the lower the legal risk. As such, a robust understanding of de-identification is important above and beyond its role in anonymisation.

These other options in legal terms involve applying measures to safeguard the data, typically under the risk-based legal ground in the GDPR known as "legitimate interests". Examples of the types of measure that may be used can be found in Annex B. These risk-based measures should generally be considered in combination with the principle of data minimisation.

To determine which option is more appropriate, we recommend that organisations:



1. Define your commercial, legal, and other requirements. What data are you using, and what do you want to do with it? What are the utility and usability constraints of your data?
2. Consider your legal options. Is anonymisation necessary, or can you achieve your goal within the scope of data protection law, utilising, for example, pseudonymisation as a safeguard to protect the data while allowing you to achieve your commercial/operational goals?

If you then decide you want to anonymise the data:



1. Decide how you will assess and document re-identification risk. You will need a robust methodology for assessing the re-identification risk in order to determine when that risk is sufficiently low for the data to be considered anonymous within the legal definition. This may involve theoretical tests as well as actual tests. You will need to document this assessment in order to meet the accountability requirement under the GDPR.
2. Select and apply controls. You will need an auditable decision making process to determine how you apply controls. This means choosing appropriate controls from the toolbox and applying them correctly. We've seen that some types of control are more effective when applied to certain types of data (e.g. generalisation works well for relatively small tables of demographic data). Some controls (e.g. perturbation or differential privacy) require you to make decisions about degree; adding more noise reduces privacy risk and, often, data utility.

Consistency and auditability are important because they can enable compliance with the GDPR's accountability requirement. A comprehensive approach to data governance will enable you to record decisions (for audit and transparency) and speed up decision making (consistent, precedent based).

3. Evaluate re-identification risk and apply further controls if necessary. You may also want to test for utility and usability for your intended purpose.
4. Audit. It may be necessary to review datasets over time to see if they're vulnerable to new risks, or to check they have been deleted as scheduled.

Data offers important new opportunities, from breakthroughs in health research to fighting financial crime, but these opportunities often lie in complex systems with context specific challenges. Genomics research could lead to life saving breakthroughs, but genome sequences pose unique and difficult privacy challenges. Financial crime is a vast and growing problem, and one where information sharing could prove vital, but concerns over privacy are a legitimate blocker.

These options may provide a way forward for both, and many other important fields, but careful consideration is needed on the specifics of each context. Privitar and Bristows are interested in exploring these areas, and others. If these are challenges you face, or you're interested in finding out more, please get in touch.



Marc Dautlich

Partner, Bristows

marc.dautlich@bristows.com



Guy Cohen

Head of Policy, Privitar

guy.cohen@privitar.com



Marcus Grazette

Europe Policy Lead, Privitar

marcus.grazette@privitar.com

Annex A - Relevant case law

This Annex provides a more detailed account of the key UK cases. However, the summaries below are not, and should not be treated as, comprehensive summaries of the cases. They highlight issues from selected cases relevant to the commentary in this paper. The cases are presented in reverse chronological order, from *Common Services Agency* in 2008 to *Peters* in 2019. We use the short titles, in italics, to refer to the cases in the body of the paper.

1. **Common Services Agency v Scottish Information Commissioner** [\[2008\] UKHL 47](#) (*Common Services Agency*)

- One of the key legal issues in anonymisation, as summarised in Section 4, is whether, when a controller puts personal data through a process to safeguard information about individuals to whom the data relates (for example, the process of ‘barnardisation’, as described in Annex B), retention of the underlying ‘raw’ dataset means that the derived dataset (the barnardised dataset) constitutes personal data in the controller’s hands.
- This case is a leading authority, but a complicated one which has not always been followed by subsequent courts, on this question.
- The Agency appealed against a decision of the Scottish Information Commissioner that “barnardised” data did not constitute “personal data” and so was required to be disclosed pursuant to a Scottish freedom of information request.
- Whether barnardised data constituted personal data in this case was a question of fact for the respondent to determine. According to Lord Hope’s leading judgement, the question was “whether the data controller, or anybody else who was in possession of the barnardised data, would be able to identify the living individual or individuals to whom the data in that form related.”
- The fact that the data controller retained access to the raw personal data from which the barnardised data was derived did not necessarily mean that the barnardised data remained personal data in their hands. If the barnardised data contained nothing that would assist in enabling a living person to be identified in combination with other information, it would not constitute personal data and could be disclosed.
- As the Commissioner had simply not considered whether barnardised data could constitute personal data, the matter was remitted to the Commissioner to consider this and make the decision afresh.

- One of the other judges, Baroness Hale, offered an alternative approach in her judgement. It was not determinative in the case. She took the position that data can be functionally anonymous, meaning that data can be anonymous from one party's perspective but identifiable from another party's perspective. In practice, the Agency is able to re-identify the de-identified data extract, because it retains the original data. However the de-identified data extract is anonymous from the perspective of the recipient, because they would not have access to the original data and would therefore be unable to re-identify the extract.

2. **Craigdale Housing Association v Scottish Information Commissioner** [\[2010\] CSIH 43](#) (*Craigdale*)

- This case confirmed the principle established in *Common Services Agency* that when considering what means are reasonably likely to be used to identify an individual, organisations should consider both the means reasonably likely to be used by the ordinary man in the street and also by a determined person with a particular reason to want to identify an individual.
- The Housing Association appealed against a decision of the Commissioner that statistics on the number of sex offenders residing in certain postal areas was personal data and should not be disclosed pursuant to a freedom of information request.
- *Craigdale* provides important pointers to the question of which attacker profiles organisations should consider when using the attack-based assessment of re-identification risk described in the motivated intruder test.

3. **Department of Health v Information Commissioner** [\[2011\] EWHC 1430](#) (*Admin*) (*Department of Health*)

- This case holds that identifiability should be assessed from the perspective of the data recipient. It supports the idea that data can be functionally anonymous and that the original source data does not need to be deleted in order to create an anonymous extract.
- Mr Justice Cranston noted the difficulty of interpreting Lord Hope's judgement in *Common Services Agency* (and the attractiveness of Baroness Hale's alternative reasoning in that case, which judicial hierarchy barred him from adopting).
- He navigated the issue by concluding that *Common Service Agency* related to the status of the data once disclosed, not when it was in the controller's hands. As such, the relevant test is whether data is identifiable by "the public to whom statistical data was disclosed".

4. Information Commissioner v Magherafelt DC, [\[2012\] UKUT 263 \(AAC\)](#) (*Magherafelt DC*)

- *Magherafelt DC* establishes the principle that, when assessing re-identification risk, the broader context should be taken into account.
- The case concerned de-identified data relating to disciplinary action taken against local council employees in a small, tight knit community. In assessing re-identification risk, the Tribunal considered the fact that council employees were well known to the local population and the council itself was a small organisation.
- The Tribunal concluded that the de-identified data was personal information on the basis that a motivated intruder, such as an investigative journalist, would have little difficulty in making the necessary enquiries which could lead to the identification of individuals subject to disciplinary proceedings.

5. Farrand v The Information Commissioner and the London Fire and Emergency Planning Authority [\[2014\] UKUT 0310 \(AAC\)](#) (*Farrand*)

- *Farrand* held that data is personal data if it can, as a whole, be related to a living individual, not whether an individual can be identified from any particular part of the data.
- The data in question was a fire investigation report, which contained several pieces of information including photographs of a flat. Mr. Farrand argued that the photographs were not personal data because they did not allow an individual to be identified.
- The Court rejected Farrand's argument that the information was not personal data because it could not itself identify any given person. It held that the photographs, taken together with other information in the report, would identify individuals to whom they related. In so doing, the Court distanced itself from the difficult reasoning relating to the status of anonymised information in the data controller's hands found in Lord Hope's judgement in *Common Services Agency*.

6. Google Inc. v Judith Vidal-Hall, Robert Hann, Marc Bradshaw v The Information Commissioner [\[2015\] EWCA Civ 311](#) (*Vidal-Hall*)

- *Vidal-Hall* raises, but does not fully answer, the question of how environmental controls should be taken into account when assessing re-identification risk.
- The relevant question was whether browser generated information (BGI) held by Google was personal data. The defendant submitted three arguments, with the first two being most relevant to anonymisation. The defendant argued that (1) BGI was, on its own, anonymous because it did not name or identify any individual and (2) BGI was segregated from other data from which an individual could be identified, such as Gmail account information.

- The court rejected the first argument on the basis that the BGI was, on its own, personal data because an individual could be singled out based on their browsing habits.
- The court also rejected the second argument, on the basis of a literal interpretation of the 1998 Act. The Act refers to information in the possession of the controller which “can be used” to identify an individual. On a literal reading, the court held that it was immaterial whether the controller intended to use or actually did use that other information to identify an individual.
- *Vidal Hall* is relevant to the question of the efficacy of use of environmental controls in seeking to restrict whether information “can be used”. *Vidal-Hall* suggests that simply storing the information separately within an organisation is not sufficient to succeed with the argument that there is no personal data in scope. It does not however settle the question whether an internal policy banning data linkage or strict access controls might fare better on this point.
- The “can be used” wording comes from the 1998 Act, so is no longer current. The current wording in the GDPR asked whether the information in question is reasonably likely to be used to identify an individual.

7. **Breyer v Germany (C-582/14), [\[2017\] 1 WLR 1569 \(Breyer\)](#)**

- This case considered whether information (a dynamic IP address) is personal data in the hands of an online service provider. The Court focused on the means likely to be used, concluding that it was not reasonable to assume that the service provider would approach any third party holding additional information. Instead, the Court held that the Directive must have been referring to only those third parties whom the service provider might approach in a reasonable manner.
- On the facts of the case, the service provider was the German government and there was a legal mechanism for it to request additional information from the internet service provider. This additional information would allow the individual user of the dynamic IP address to be identified.
- The court follows the Advocate General’s opinion in excluding criminality from its assessment of identifiability. The Advocate General concluded that a means of identification is not reasonably likely to be used “if the identification of the data subject was prohibited by law or practically impossible on account of the fact that it requires a disproportionate effort in terms of time, cost and man-power, so that the risk of identification appears in reality to be insignificant.”
- The court’s approach in *Breyer* is more pragmatic than the A29WP’s absolutist approach in Opinion 05/2014. In accepting the ‘insignificant-in-reality’ threshold, *Breyer* diverges from Opinion 05/2014 which suggests that re-identification must be ‘impossible’.

8. Peter Nowak v Data Protection Commissioner (C-434/16) ([ECLI:EU:C:2017:994](#)) (Nowak)

- The CJEU confirmed that identifiability is considered in the abstract, rather than from the perspective of a particular party (e.g. the data recipient) and that a single party does not need to hold all of the information necessary to identify an individual.
- The CJEU was asked to determine whether “the written answers submitted by a candidate at a professional examination and any examiner’s comments with response to those answers” constituted personal data.
- The Irish Data Protection Commissioner argued that the script was not personal data because the examiner marking it could not identify a candidate. The court disagreed. It held that it was sufficient that anyone, in this case the exam board, could identify the candidate.

9. Information Commissioner v Miller [\[2018\] UKUT 229 \(AAC\)](#) (Miller)

- This case sheds light on the question of whether it is easier for some types of data (including data judged to be of only ‘historic’ interest) to pass the legal threshold to be considered anonymous. (Yes, it is). Whether the data pass this test depends on whether the data would be likely to attract the interest of a ‘motivated’ intruder, that is, one more skilled than ‘the man in the street’.
- The appeal concerned whether historical homelessness data for a particular local authority area constituted personal data exempt from disclosure in response to a freedom of information request. The Tribunal found that the data in question, both because of its nature and its age, was unlikely to motivate an intruder to attempt to re-identify any individual to whom the data related.
- The Tribunal provided two examples of motivated intruder profiles: (1) a member of the public and (2) an intruder “with investigative skills, such as a journalist”.
- The Tribunal noted that “information in the spreadsheets is not such as is likely to attract those with investigative skills, such as a journalist, to attempt to identify individuals” and “even if, which is unlikely, there may be some interest in [individuals to whom the data relates] at or close to the time, I can see no basis for thinking that it would be of interest to anyone several years later.”
- The Tribunal then considered the threshold for identifiability. The Tribunal concluded that, based on the data in question, “the chance of a member of the public being able to identify [an individual] from the data is so remote as to be negligible”.

10. John Peters v (1) the Information Commissioner and (2) the University of Bristol [2019] UKFTT 2018 0142 (Peters)

- *Peters* summarises the characteristics of the motivated intruder test as follows:
 - The intruder starts without any prior knowledge,
 - He or she will have a particular reason to identify the subjects of anonymised data and will take all reasonable steps to do so, will be determined and competent, will have access to resources such as the internet, libraries and public documents but will not resort to criminality,
 - He or she is similar to an investigative journalist but will not have specialist equipment,
 - Non-recorded personal knowledge about the data subject (e.g. that of a doctor or family member) can be relevant but there must be a plausible and reasonable basis for believing that such knowledge presents a significant re identification risk,
 - The intruder must have a likelihood of success, and
 - Educated guesswork does not suffice.
- The Tribunal also observed that there should be an assessment of whether there is a realistic possibility of a motivated intruder being interested in the data in a given case, though it did not need to make a ruling on that point for the purposes of its decision and accordingly left this important point open.

Annex B - Common De-identification Techniques

Data Transformations

Data transformations aim to mask values or to make data less precise. The most commonly used masking techniques are:

1. Deletion (full redaction). This is an important aspect of data minimisation, if a value is not necessary for the analysis it should be deleted by default. Deletion can mean removing the records entirely (e.g., removing an attribute/column), or replacing all values with a constant value such as “XXXXX” or 0.
2. Clipping (partial redaction), where a value is partially deleted. This is commonly applied to credit card numbers, where only the last four digits are retained. This can also have the result of generalising the data (see point 8 below).
3. Tokenisation, where a value is replaced with a randomly generated value, a token. Tokens may need to conform to a specific format, e.g., a specific length, to ensure that tokenised data remains compatible with other processing that may be applied to it. For example, UK National Insurance numbers conform to a specific format (two prefix letters, six digits and one suffix letter) and follow specific rules (e.g. the second prefix letter is never an ‘O’).⁷ Generating a token with the correct formatting allows tokenised data to pass a validation test.
4. Hashing, where a function is applied to the value to produce a fixed length output known as a “hash”. The function is one-way, so the hash cannot be converted back to the original value. This is a common technique but has been shown to be vulnerable to attack.⁸ Hashing has many variations, including salted hashing where a random string, a “salt” is added to the value before it is hashed.
5. Substitution, where a value is replaced by another value from a predefined list. This can provide a form of generalization, where the substituted value is less precise than the original and many values map to it. For example, the values “Westminster” and “Lambeth” could both be substituted for the more general value “London”.
6. Field level encryption, where a value is encrypted to produce an output ciphertext derived from the input value and a cipher key.

⁷ HMRC, National Insurance Manual. Accessed August 2020

⁸ Vijay Pandurangan, On Taxis and Rainbow Tables: Lessons for researchers and governments from NYC’s improperly anonymized taxi logs. Accessed June 2021.

7. Perturbation, where random noise is added to a value, e.g., transactions could be perturbed by any full unit value in a range, so an input value of \$182 could be perturbed by +/- \$5 to generate an output value in the range \$177-\$187.
8. Barnardisation is a type of perturbation applied to tables of counts for statistical disclosure control purposes. It involves randomly adding or subtracting 1 from some cells in the table.
9. Generalisation, where a value is made less precise. The specific technique varies depending on the type of data:
 - Binning. This most commonly involves transforming a specific value into a range, e.g., £182 transformed to a range £180 - £190, or a midpoint £182 → £185.
 - Rounding. E.g., to the nearest 100, so £182 → £200.
 - Clipping some types of values. E.g., clipping a postcode to just the first three digits, or clipping a time to give just the hour, not the minutes or seconds, so 09.06.55 → 09.

Environmental Controls

Environmental controls are much broader. We can group them according to the type of behaviour that they are intended to prevent:

- Prevent unauthorised access - e.g. authenticated access, data recipient vetting, data in standalone server, physical access to machines/workstations
- Prevent unauthorised processing - e.g. standardised data sharing agreements, data processing within restricted secure environment
- Prevent data use beyond agreed purpose(s) - e.g. data sharing agreements to require destruction of data after a set period or restricts potential use
- Deter malicious behaviours - e.g. data recipient vetting, logging or monitoring interactions with the data, flagging warnings following modification or disappearance of data
- Prevent accidental data misuse - e.g. staff training, auditing of data handling processes and systems, clearly defined procedures

About

Bristows

Bristows is a European headquartered law firm for litigation, transactions and advice. We have one of the largest and best known data protection teams in Europe. We specialise in providing practical and proportionate data protection advice, from compliance programmes to cyber incidents, for organisations large and small.

[Visit the website.](#)



Privitar is the leader in modern data provisioning. We empower organizations to use data safely, quickly and at scale. Our clients use the Privitar Data Provisioning Platform to share data, unlock insights, keep data safe and support regulatory compliance. Our platform includes state-of-the-art privacy enhancing technologies, and our experts help customers use them effectively.

[Visit the website.](#)

About the authors



Marc Dautlich
Partner, Bristows

marc.dautlich@bristows.com

Marc is a technology law specialist with over 25 years' experience gained in private practice and as inhouse legal counsel. He advises clients in the technology, financial services and life sciences sectors on technology, data protection and cyber risk matters and is a member of the International Technology Law Association, the International Association of Privacy Professionals, the Data Protection Finance Group, the Cyber-Security Information Sharing Partnership, and TechUK.



Guy Cohen
Head of Policy, Privitar

guy.cohen@privitar.com

Guy joined Privitar in 2016, prior to which he worked in the UK Civil Service, in the Department of Health, the Cabinet Office and HMRC. Guy has been a fellow at Cambridge University's Centre of Science and Policy, a member of the Royal Society Privacy Enhancing Technologies Working Group, and is the technical editor for the IEEE Data Privacy Process Standard.



Marcus Grazette
Europe Policy Lead, Privitar

marcus.grazette@privitar.com

Marcus is a European public policy specialist. He joined Privitar in 2019 after an eleven-year diplomatic career with the UK's foreign ministry, including a secondment to EY as a consultant. He studied at the *École nationale d'administration*, and holds an MA in European public policy from *l'université Paris 1 Panthéon-Sorbonne*.

Bristows  **PRIVITAR**