

Transparent Government, Not Transparent Citizens: A Report on Privacy and Transparency for the Cabinet Office

Kieron O'Hara

*Intelligence, Agents, Multimedia Group
School of Electronics and Computer Science
University of Southampton
Highfield
Southampton SO17 1BJ
United Kingdom
kmo@ecs.soton.ac.uk*

Executive summary and recommendations

In December 2010, I was asked by the Minister for the Cabinet Office to conduct a review about the issues for privacy that were raised by the Coalition government's transparency programme. During the review period, experts in government, civil society activists, academics and many others were consulted to try to reconcile the desire for open government with the privacy of individual citizens (who may be data subjects in datasets about government activity). Those who were kind enough to help the review are acknowledged at the end of the report.

The review reached the following conclusions.

- Privacy is extremely important to transparency. The political legitimacy of a transparency programme will depend crucially on its ability to retain public confidence. Privacy protection should therefore be embedded in any transparency programme, rather than bolted on as an afterthought.
- Privacy and transparency are compatible, as long as the former is carefully protected and considered at every stage.
- Under the current transparency regime, in which public data is specifically understood not to include personal data, most data releases will not raise privacy concerns. However, some will, especially as we move toward a more demand-driven scheme.
- Discussion about deanonymisation has been driven largely by legal considerations, with a consequent neglect of the input of the technical community.
- There are no complete legal or technical fixes to the deanonymisation problem. We should continue to anonymise sensitive data, being initially cautious about releasing such data under the Open Government Licence while we continue to take steps to manage and research the risks of deanonymisation. Further investigation to determine the level of risk would be very welcome.

- There should be a focus on procedures to output an auditable debate trail. Transparency about transparency – metatransparency – is essential for preserving trust and confidence.

Fourteen recommendations are made which are intended to implement these conclusions without making too strong a claim on resources.

1. **Represent privacy interests on the Transparency Board.**
2. **Use disclosure, query and access controls selectively.**
3. **Include the technical paradigm.**
4. **Move toward a demand-driven regime.**
5. **Create a data asset register.**
6. **Create sector transparency panels.**
7. **A procedure for pre-release screening of data to ensure respect for privacy.**
8. **Extend the research base and maintain an accurate threat model.**
9. **Create a guidance product to disseminate best practice and current research in transparency.**
10. **Keep the efficacy of control in the new paradigm under review.**
11. **Maintain existing procedures for identifying harms and remedies.**
12. **Use data.gov.uk to raise awareness of data protection responsibilities.**
13. **Investigate the Vulnerability of Anonymised Databases.**
14. **Be transparent about the use of anonymisation techniques**

The grounds for these conclusions and recommendations are given in the body of the report, and the recommendations elaborated in detail in the final section.

1 Introduction

Transparency as practised in the United Kingdom is a very new and innovative phenomenon. As recently as 2007, a trio of political scientists from Harvard's Transparency Policy Project wrote one of the most important studies of transparency (Fung et al 2007), they focused on the mandatory publication of data by (usually) private sector outfits for particular purposes – not the voluntary publication of lots of data about anything and everything, available to anybody and everybody. This latter type of transparency (they called it 3rd generation transparency) they saw as “a glimpse of a technology-enabled future”, but erroneously imagined that government, though a coordinator of efforts, would not be a provider of data (Fung et al 2007, 151-169).

Similarly, as transparency is an innovation that came, if not out of the blue, at least very unheralded, the potential threat to privacy has not been considered and theorised to any great extent. In September 2008, a special edition of *Scientific American* featured 12 articles on digital privacy by some of the finest commentators in the field, including Simson Garfinkel, Whitfield Diffie and Daniel Solove. Less than three years ago – yet transparency of government data was not mentioned once in the issue. One of the leading introductory textbooks on US privacy law, dating from 2008, does not include transparency in its closing survey of future challenges (Soma & Rynerson 2008, 292-341)

Yet we do need to think about these issues. A good government will have a range of interactions with its citizens, but it is essential that transparency of government does not lead to exposure of the citizen. This report will consider ways to prevent this from happening.

I shall argue that a proper concern for privacy is not incompatible with transparency. The proceedings of government can, and should, be open to scrutiny without compromising citizens, or sacrificing their seclusion, or preventing them from keeping control of their self-presentation. I am optimistic that those pushing forward the transparency agenda in the United Kingdom are sensitive to privacy concerns, and will act accordingly.

Indeed, not only are privacy and transparency compatible, privacy is *a necessary condition* for a successful transparency programme. Transparency requires public confidence, and one way to ensure that is to reassure the public that its privacy is a central concern whose protection is embedded in decision-making processes. This reassurance will only happen if the transparency programme is itself transparent, so that discussions and debates are open to inspection.

In this report, I shall set out recommendations that I believe will allow the integration of privacy protection with transparency, and help preserve public confidence. My first task is introductory. Section 1.1 will set out the basic ideas and practices of the UK's transparency programme, to specify the context of this review. In section 1.1 I shall discuss the terms of reference of the review, and explain my interpretation of it – in particular, why I feel that privacy is the relevant and most important driver here, as opposed to the legalistic idea of data protection. Section 1.3 will then outline my approach, while section 1.4 will explain the report's overall structure.

1.1 *The practice of transparency in the United Kingdom*

The transparency programme of the UK government is a specific example of the transparency ideology in action. I will discuss the broader ideology in section 2.1 below, but in this section I shall set out the immediate context which may be of value to the reader.

1.1.1 The Coalition Agreement

Transparency is an important part of the Coalition government's political agenda. The Coalition Agreement (Cabinet Office 2010) states:

The government believes that we need to throw open the doors of public bodies, to enable the public to hold politicians and public bodies to account. We also recognise that this will help to deliver better value for money in public spending, and help us achieve our aim of cutting the record deficit. Setting government data free will bring significant economic benefits by enabling businesses and non-profit organisations to build innovative applications and websites.

On this basis, the agreement makes a number of specific commitments, augmented by the Prime Minister's letters of 29th May, 2010 and 7th July, 2011 (Cameron 2010, 2011).

Transparency is clearly central to the government's plans.

1.1.2 The government's aims

The government has set itself the target of making the UK the most transparent and accountable government in the world. By doing this, it hopes to achieve the following aims.

- Making government more accountable and approachable, by moving from administrative accountability to more direct democratic accountability, enabling citizens to hold the government to account.
- Creating better value for money by providing an insight into how money is spent, encouraging departments to improve controls on spending and reduce their costs.
- Stimulating growth by enabling businesses to develop innovative information-based products and applications using public data.
- Reforming public services by:
 - Providing choice and improving public sector outcomes, by giving citizens the information they need to make informed decisions about the public services they use, and giving providers the incentives they need to improve the quality of their services and to develop new innovative services.
 - Opening up public sector contracts, giving companies, social enterprises, charities and employee-owned cooperatives the opportunity to compete to offer high quality services by providing access to public sector contract and procurement data.

1.1.3 Institutional structures

To implement these plans, the Cabinet Office contains a Transparency Team tasked with delivery of the transparency programme. An advisory body, the Public Sector Transparency Board, chaired by the Minister for the Cabinet Office, is charged with driving the policy. The Board has released a set of Public Data Principles which are intended to ensure that data releases are timely, valuable and reusable (Transparency Board 2010). A Local Public Data Panel plays a similar role to the Transparency Board with respect to data from local government. A website, data.gov.uk, is intended to act as an aggregator for public data releases in open and standardised formats.

In March, 2011, it was announced that Tim Kelsey of McKinsey's would be an advisor to the government, supporting the Government in shaping its transparency agenda over a period of at least six months.

A new UK Open Government Licence (OGL – <http://www.nationalarchives.gov.uk/doc/open-government-licence/>) was launched in September 2010 as a simple and straightforward set of rules to enable people to re-use government data in any way they want. The OGL implements the commitment not only to publish the data but to allow everyone to use it freely, helping to create a new era of social entrepreneurs. The licence is available in machine readable form, flexible and works in parallel with other internationally-recognised licensing models such as Creative Commons. The new UK licence does not require users to register or formally apply for permission to reuse data.

1.1.4 Government agencies and third party suppliers

The selection of data for release can be made on a number of grounds, including ease of publication, likely value for the public and so on. Currently the Transparency Board takes a lead in pushing for particular datasets to be published.

It should be noted, however, that practical considerations are not the only ones driving decisions about what and what not to publish. Most government departments have discretion via common law powers over what they can and cannot publish – in other words, where they are mandated in law neither to restrict nor provide access to data, they can make a choice. Forthcoming right to data legislation is intended to promote data releases and give the public greater rights to ask for data (it also specifies that data should be released where feasible in reusable form).

Other statutory bodies whose managers lack the powers of ministers are somewhat more restricted in what they can do, and their decision-making powers in this space are determined by legislation. Some bodies have very strict requirements of confidentiality laid upon them by the relevant acts of Parliament. An example here would be the 2005 Commissioners for Revenue and Customs Act, which is intended to protect taxpayer confidentiality within Her Majesty's Revenue and Customs.

In particular, it is worth noting that, although most such legislation would have been enacted during periods of government when transparency was not valued, where duties of confidentiality have been created by legislation there is usually good reason. In considering the privacy implications of data releases, the reasoning behind such legislation, where it exists, should be a factor in debate.

With regard to non-governmental organisations supplying services to government, who may generate valuable data as a by-product of a government contract, one would not want to see a smaller commitment to transparency. It would be unfortunate indeed

if a more efficient government that outsourced more functions became less transparent because it generated less data directly.

In the case of both non-departmental government agencies and non-governmental organisations, there is already a set of scrutinising principles in place to determine whether a release of information is justified (and to what extent it can/should be redacted) under Freedom of Information legislation. Given the soundness of these principles, there seems to be little reason why transparency should not apply to these non-departmental bodies, by following the FoI principles but with a proactive publication strategy to address public demand, rather than a reactive publication strategy driven by orders from the government.

1.1.5 Demand and supply, and the right to data

Currently, releases of data are driven largely by the Cabinet Office, the Transparency Board and individual departments and agencies. Forthcoming right to data legislation, as noted, will give citizens a say in what information and data are provided.

Hence it is fair to say that the transparency programme of the UK government at the moment is largely a top-down process. It is an aim of the Transparency Board and the Transparency Team to move toward a situation where demand for data was easy to register. In such a situation, the transparency programme would become a more bottom-up, demand-driven process. That would be very much more in line with the underlying philosophy of transparency.

The Cabinet Office is currently working to amend the Freedom of Information Act 2000 to ensure that all datasets realised through FoI must be available for reuse and in a reusable format, available to everyone and able to be exploited for social and commercial purposes. At the time of writing these legislative changes are included in the Protection of Freedoms Bill. Until the Bill is enacted, the advice is that public requests to departments for the release of government datasets should be handled in line with a presumption in favour of transparency, with all published data licensed for free reuse including commercial reuse. Since these are data which departments already have, or should have, this is not expected to involve significant costs or new IT systems.

1.2 Terms of reference of the review

The terms of reference that governed this review were set out in a letter from the Minister for the Cabinet Office. In particular, I shall focus on the sections of the terms that specify the purpose and remit of the review.

1.2.1 Content of the terms of reference

The terms of reference of the review were as follows:

Purpose

The Minister for the Cabinet Office has established a review of the impact of Transparency on Privacy. The Review will ensure that as the Government develops its transparency agenda, it continues to uphold high standards of personal privacy.

The Review will support officials and Ministers in ensuring that on-going releases of data are done in a way that provides maximum transparency of data consistent with the appropriate data protection safeguards.

Remit

The Review will:

- Support the Government in striking the right balance between transparency and data protection safeguards, and between the interests of wider society and the interests of the individual or corporate body.
- Identify the nature of the risk to privacy of the individual posed by transparency of public data, in particular the potential for 'jigsaw' identification.
- Advise the Government on practical approaches to take.

1.2.2 Legal concepts of 'privacy' and 'personal data'

Privacy is generally discussed in primarily legal terms, using instruments developed over many years including the relevant article (article 8) of the European Convention on Human Rights. Article 8 states:

Article 8 – Right to respect for private and family life

1. Everyone has the right to respect for his private and family life, his home and his correspondence.
2. There shall be no interference by a public authority with the exercise of this right except such as is in accordance with the law and is necessary in a democratic society in the interests of national security, public safety or the economic well-being of the country, for the prevention of disorder or crime, for the protection of health or morals, or for the protection of the rights and freedoms of others.

Other than this, there is no independent privacy tort (i.e. a wrong caused by a failure to perform a civic duty to respect privacy) in UK law (I shall discuss the issues of privacy harms and remedies in section 3.3). Note for now that Article 8 has a number of get-out clauses which provide grounds for transparency activists to contest a privacy ruling. Indeed, these clauses are a useful checklist for transparency activists – if a privacy-threatening data release does not help prevent disorder or crime, or protect rights and freedoms, or help national security or economic well-being, why go through with it?

The other important instrument in this area is the Data Protection Act 1998, based on the EU Data Protection Directive 95/46. The Act, like all data protection legislation, is not specifically intended to protect privacy, but rather to balance the interests of the subjects of data with the interests of data users (for a review of data protection legislation, see Walden 2007).

The Directive defines personal data as follows.

- (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.

This appears to be clear, and indeed was clearer in the days when linking data across applications was difficult. The lack of ability to transfer data easily from host to host

created a kind of practical obscurity that the Directive – already becoming out of date before it was published – trades on. This bygone nature of the Directive may be taken as somewhat worrying, given that, in the absence of a specific privacy tort in UK law, the discourse of data protection is “assimilating privacy questions” and so, to an extent at least, is being used by judges as a means of protecting privacy (Wacks 2006, 173).

1.2.3 Issues with data protection

The Data Protection Directive, and the associated Act of Parliament which implements it, appear to be somewhat unloved by both privacy advocates and those who champion data-sharing as a means to efficiency and effectiveness. This might seem odd given the opposite pull of these two positions, but is explained by the lack of clarity of the Directive which renders much uncertain, particularly when considered in the light of the development of new technology and the evolution of social attitudes to technology that have taken place in the last twenty years.

The form of an EU directive is that the substantive provisions are given in *articles*, which must be implemented by national governments within a reasonable timeframe. Interpretations of the directive, including its aims, objectives and the background context understood by its authors, are given in a series of *recitals*, which do not have to be implemented in national legislation. The Data Protection Directive has been implemented unevenly across the EU, leading to many uncertainties (Korff 2003).

To take an obvious and oft-criticised example, an identifiable person is defined as “one who can be identified”. We can adopt the gloss of ‘identifiable’ produced by the European Commission’s advisory Article 29 Working Party here.

In general terms, a natural person can be considered as “identified” when, within a group of persons, he or she is “distinguished” from all other members of the group. Accordingly, the natural person is “identifiable” when, although the person has not been identified yet, it is possible to do it (that is the meaning of the suffix “-able”). ...

Identification is normally achieved through particular pieces of information which we may call “identifiers” and which hold a particularly privileged and close relationship with the particular individual. Examples are outward signs of the appearance of this person, like height, hair colour, clothing, etc... or a quality of the person which cannot be immediately perceived, like a profession, a function, a name etc. The Directive mentions those “identifiers” in the definition of “personal data” in Article 2 [quoted above, section 1.2.2]. (Article 29 Data Protection Working Party 2007, 12)

This raises the natural question: ‘identifiable by whom?’ An act implementing the Directive could be very strong if the answer to this question is ‘by anybody’, or relatively weak if the answer is ‘by the data controller’.

The question of ‘identifiable by whom?’ is discussed in the Directive’s Recital 26.

(26) Whereas the principles of protection must apply to any information concerning an identified or identifiable person; whereas, to determine whether a person is identifiable, account should be taken of all the means likely reasonably to be used either by the controller or by any other person to identify the said person; whereas the principles of protection shall not apply to data

rendered anonymous in such a way that the data subject is no longer identifiable
...

This is quite strong, suggesting that if *anyone* can identify someone from the data it is personal data. However, the upshot of this recital was not implemented in the Data Protection Act, which instead puts the onus on the *controller*:

“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;

If we examine Recital 26 further, it raises more questions. All means “likely reasonably to be used” – this is an extraordinarily vague phrase. As shall be discussed later, there are methods for deanonymising data that are very powerful, and although they are probably (at the time of writing) beyond the means of most individuals, their cost is balanced by the potential for exploitation of a large quantity of data in a dataset.

Similarly, the principles of protection are not applied to anonymised data – but anonymised in the Recital simply means that “the data subject is no longer identifiable.” Once more, this raises the question: ‘identifiable by whom?’ Recursively, one would imagine that the Recital intends the same wide interpretation of the question as it itself applied to the Directive, which means the same issues arise. In a world in which there are very powerful techniques for deanonymising data, is the upshot of the Recital that data sharing and transfer should virtually cease (as, for example, the legal scholar Paul Ohm has argued is a very reasonable interpretation of the Directive as a whole – Ohm 2010)?

Furthermore, how effective do the means have to be to suggest that a person should be identifiable? Suppose I have a deanonymisation method that enables me to identify a person in a given dataset with a probability of 0.1% – I can identify one person in a thousand. Then any given person in the dataset is very unlikely to be identified by my method, but if I had access to a dataset with sensitive information about 25m individuals (such as the Child Benefit database lost by HMRC in 2007), then I could be reasonably sure that I could get access to sensitive information about 25,000 individuals – not a bad haul. For any individual person, it could be said that the chances of his being identified by these means are very small. Nevertheless, it is clear that many individuals are likely to be identified by the method.

Of course, these and other interpretative questions have been thrashed out in various courts over the years, but grey areas remain both at the European level and at the level of national law (Korff 2003). The Directive is also being revised at the time of writing (http://ec.europa.eu/justice/policies/privacy/review/index_en.htm). We need to ask whether data can be linked to an actual person, at what cost, with what effort, for what purpose, with what likelihood, before we can even start to make a judgement. Answering these questions is non-trivial, and demands not only legal but economic and technical knowledge.

1.2.4 Interpretation of the notion of 'privacy' in the context of the terms of reference of this review

Clearly, the legal definition of data protection, and its relation to privacy, is extremely important in this space. It has come under criticism from privacy activists who believe that its protection is too scanty, from data users who believe that its provisions undermine the fruitful exploitation of data for the public good, and technologists who believe that it is based on an outdated understanding of the technical capabilities of those who wish to undermine personal privacy. Some lawyers have argued that the legal definition is incoherent (Wacks 2006, 175-181). Nonetheless, it should go without saying that, however inadequate the law is, the government should not wilfully break it. Furthermore, the law represents an attempt to codify social norms pertaining to privacy, and so has to be taken extremely seriously as a guide, if a flawed one, to what is acceptable and what not acceptable.

Nevertheless, the most cursory of examinations of the terms of reference will show that the focus on data protection is not on its own adequate for this review. Firstly, the review is tasked with advising the government on practical approaches to ensuring that ongoing releases of data can continue without compromising privacy. That certainly entails that the government does not break the law, but also entails that public confidence in the transparency programme is maintained. If there was a perceived problem with a data release that affected many thousands or millions of people, then, even if the government had adhered strictly to the letter of the law, the transparency programme would lose much of its political legitimacy. Equally, on the other side, worries about data protection could lead to unwarranted risk aversion about transparency, which will have its own chilling effect on the use and release of data in the public domain.

Hence this review must also take into account public perceptions of privacy and private life. Although legal definitions have attempted to codify these, public perceptions change over time, and are rarely informed by the state of the law at any one time. Technological innovation is a particularly speedy driver of public perceptions – consider how behaviour on social networking sites has surprised many observers. The public is not a homogenous group of people, and its perceptions cannot simply be enumerated, but equally there has been important work both theoretical and survey-based on how the public views its privacy and invasions of that privacy, which this review will also take into account (cf. e.g. Bradwell 2010, Coles-Kemp et al 2010). It should be said at the outset that there is evidence that the public is more relaxed about privacy, particularly as a value to be traded off against good or improved public service, than experts and privacy campaigners often realise (cf. Kelsey 2009, Bradwell 2010).

So, to take one example of how one should reason in this space, some transparency activists argue, with logic, that if data *could* be released reactively under the terms of the Freedom of Information Act (FoIA), then they *should* be released proactively under a transparency programme, as there can be no *legal* impediment. This is a pleasingly elegant formulation – yet one that I believe should be resisted, because *public* perceptions of privacy may not be exactly expressed by FoIA. The activists' principle could result in a perfectly legal release of data that causes a loss of confidence in transparency.

The second reason to expand the inquiry beyond data protection is that the terms of reference specifically mention the risks of jigsaw identification. This is a technical

concept describing the ability or otherwise of an adversary to reidentify or to deanonymise anonymised data, with the help of background information and processing power (see section 4 for detailed discussion of this). Because of the technical element of this practice, it is also essential for this review to take into account technical definitions and concepts of privacy from mathematics and computer science.

Hence, the review will not focus directly on the legal definition of privacy, but instead will try to broker between the legal definition, the technical definition and public perceptions. The success of the transparency programme will depend on the government not breaking the law *and* on not losing the trust of the citizen, and to do that it will need, among other things, to avoid releasing information in a form which could lead to widespread compromises of privacy by adversaries *even if the releases were legal and the possibility of compromise was initially unanticipated or discounted by the public*. Hence all three paradigms – the legal, the public perception and the technical – are essential.

1.2.5 Interpretation of the notion of 'transparency'

'Transparency' covers a number of different styles of data release, including: the sharing of possibly sensitive personal data for disinterested research; the sale of data to companies or other bodies; the mandatory publication of data on certain matters (e.g. company accounts, or energy levels of electrical goods); and a Freedom of Information regime within government. *These are not the focus of this review.*

Its focus is the release of datasets not to individuals, but to everyone, in reusable form, with few restrictions of use (e.g. under the Open Government Licence), via an accessible infrastructure (such as the World Wide Web), under the administrative infrastructure described in section 1.1.

Some of the conclusions of this review could be different if dealing with different ideas of transparency. The review applies only to the type of unconditional transparency sketched in the previous paragraph.

1.2.6 Privacy, the individual and the corporation

Given this line of thought, an immediate issue presents itself with respect to corporate bodies. Of course, corporate bodies have deep and important issues with respect to transparency – for example confidentiality, copyright and intellectual property. Nevertheless, I do not believe that these issues are best dealt with under the rubric of 'privacy', however analogous they may be. They have legal bases and, unlike the privacy of individuals, are best dealt with using legal procedures and reasoning.

To take one example, there are currently in place guidelines concerning redactions from contracts to be published online based upon principles previously developed in the context of FoIA. The guidelines covering the publication of central government contracts, for instance, state:

Certain redactions may be required prior to publication in order to protect certain types of information which may be considered exempt from publication. Redactions of contractual text are permitted in line with the exemptions set out by the Freedom of Information Act. This is also the approach being taken for the requirement to publish items of central government spending over £25,000. The Freedom of Information Act contains 23 grounds for possible exemptions. For example, these exemptions may include information in relation to national

security, commercial confidentiality and the protection of personal data as permitted by the Freedom of Information Act. (Cabinet Office 2011a, 9)

Also see (Cabinet Office 2011b).

The use of a legal framework already informed by FoIA is a logical step. If information or data held by a corporation can be freed under the terms of FoIA, then it seems reasonable to say that it should be freed under the transparency programme, in the corporate context where public confidence in transparency is unlikely to be threatened.

Nevertheless, the procedures and recommendations I shall set out in section 5.3 could be adapted to the corporate case if this was felt necessary.

1.3 Approach of this review

Given the above interpretation of the terms of reference, in this section I shall sketch the approach that I have taken, particularly focusing on the properties of the solutions I shall be putting forward.

1.3.1 Broad principles

The transparency regime in the United Kingdom is a new phenomenon whose institutions are evolving extremely quickly (section 1.1). The Transparency Board and the Transparency Team of the Cabinet Office have no statutory powers to release data; neither do other agencies, such as the Home Office, for example, which had to work in partnership with police forces and other agencies to deliver the Prime Minister's commitment that the citizen should be able to see the level of crime in his or her street by early 2011. Hence the transparency programme has been until now realised by an *ad hoc* combination of exhortation, pressure, expenditure of political capital, cajoling and reasoned debate.

It follows that it would not be sensible to prescribe a particular set of institutions and relationships, be they ever so brilliantly devised. This review will instead try to describe *the debates and arguments that should take place, together with the evidence that should be amassed, that will enable policymakers to determine the extent of risk associated with a particular proposal to release data in one or another form.*

As a result, in its fourteen recommendations this review will set out the broad principles of an approach to the consideration of the privacy issues in this space, rather than detailed specifications of institutions or procedures.

1.3.2 Case-by-case reasoning

Furthermore, it is clear that the transparency programme is blazing a trail in a number of innovative ways, changing our political assumptions. The attitude to transparency of the public, of politicians, of the media and of public servants is likely to evolve unpredictably over the next few years. There will be unintended and unanticipated consequences.

Hence setting out strict principles, or an exact institutional structure, would be counterproductive in another way – any set of ideas to deal with the privacy issues that transparency will create will need to adapt to new circumstances, new demands and changing attitudes. What is possible now may not be possible in 2015, and vice versa.

There is divergence across the various areas of policy – health, transport, education, criminal justice, etc. Any particular release of data will have its own set of properties and constraints that could vary along any or all of the following dimensions:

- Public expectations.
- Legal barriers to the release of data.
- The availability and richness of background information that would be valuable to a person who wished to deanonymise personal data.
- The sensitivity of the public to privacy breaches (which itself often differs between the generations, between the sexes or between people of different educational attainments).
- The nature of the potential harms.
- The value to the public of a data release.

Hence it is important, if possible, to treat each dataset to be released as having unique characteristics based on its specific context, and deal with the privacy issues on a case-by-case basis. In this review, the broad principles and procedures I shall set out will facilitate case-by-case treatment without placing too heavy a bureaucratic burden on the transparency programme.

1.3.3 Balance

'Balance' is an important concept here. As I have argued elsewhere (O'Hara 2010), it is incorrect to assume that privacy is of value primarily to the individual, while the interests of society are served by eroding privacy. Privacy is a public good, essential for the successful functioning of a democratic society (Rössler 2005, Solove 2008, Raab forthcoming), and so the balance between privacy and transparency (which is another public good) cannot simply be expressed as a balance between the interests of the individual and those of the community.

To say this is to assume there is a linear interpretation, a zero-sum game in which a successful data release will of necessity invade privacy, and a successful defence of privacy will of necessity prevent data being used effectively for public benefit. These propositions are false: on many occasions privacy and transparency will push in the same direction.

To take one obvious type of example, the success of a transparency programme depends on public trust, which is more likely to be preserved if the public feels that those in charge of the programme respect its privacy concerns. In a second type of example, a transparency programme will be able to furnish a full and officially-sanctioned record of events, which by stating the full facts, avoiding both partiality and sensationalism, will reduce the incentives to spread misleading and decontextualized accounts which could be far more damaging even if true. Consider the publication of accurate court data: that could surely help by setting out the complete official record, recording the 'not guilty' verdicts alongside arrests, and successful appeals alongside convictions. Other media, reflected in search engines, may instead concentrate on the sensational, to create a true but misleading narrative.

Nevertheless, it has to be admitted that privacy is only one of several kinds of public good, which may well on occasion clash. When such circumstances do arise, the language of balance tends to skew the debate, particularly as privacy is often cast as

primarily an interest of the individual and a cost to the community. In one of the earliest uses of the word in *Troilus and Cressida*, Achilles, sulking in his tent, grows "Of this my privacy/I have strong reasons." The wily Ulysses replies "But 'gainst your privacy/the reasons are more potent and heroic." Shakespeare seems to be hinting that the public duties of Achilles should outweigh his private motivation.

The problem then arises that the interests of the individual can hardly be expected to be treated with equal weight to the interests of the entire community (an asymmetry that has worried a number of commentators, such as Raab 1999). So embattled is the right to privacy that balance will always be extremely difficult to achieve.

To resolve the issue we should note that 'balance' is both a noun and a verb. It can denote a state wherein goods are distributed equally and justly. However, such a state is unlikely to be easily achieved in this area. Given a potential data release, a decision must be made as to whether the release goes ahead or not; there is no balance between release or retention.

A more useful interpretation is of 'balance' as a verb denoting a process of comparison and just treatment of competing goods and interests. In particular, the status of privacy as a basic human right means that any such process must ensure that privacy is preserved. This interpretation brings other more useful notions into play, including proportionality, necessity and the public good. In a practical situation, this interpretation would introduce such questions as, for example, whether it would be possible to achieve the same or comparable effects without a release of potentially identifying data.

Hence this review will treat its balance requirement as follows: *to determine the maximum level of transparency consistent with an acceptable level of privacy in British democracy.*

1.3.4 Data and information

One more important terminological clarification concerns the well-known distinction between *information* and *data*. This distinction is more often gestured towards than defined rigorously, but broadly speaking data are at a lower level of abstraction than information. Information is data interpreted for some audience in some way, data presented in order to maximise their utility in a context. Data themselves might be a set of numerical values of some parameters, or (in the world of linked data) a set of triples in the Resource Description Framework (RDF). The crime data released monthly are data, while the maps and functionality of the police.uk website are information.

A government could be transparent with respect to information or to data. Given the interests and activities of the Transparency Board and the Transparency Team of the Cabinet Office who have commissioned this review, and given the focus in the terms of reference on data, I shall concentrate on *data* transparency – the release of datasets which can be turned into information (i.e. applied to a task) by anyone who downloads the data from data.gov.uk or some other outlet – rather than *information* transparency in this review. Many of the arguments are perfectly general over all levels of transparency, but if there is a doubt, the context is *data* transparency.

I shall restrict my comments to transparency with respect to data generated by or on behalf of public sector organisations. I shall call them 'government data' and 'public data' interchangeably. I do not consider the possibility of private sector organisations,

whether companies or non-profit organisations, being transparent in this sense *except* where they are providing services for government, competing for government tenders and so on.

1.3.5 The review will not recommend new legal or technological instruments

The aim of this review will therefore be to suggest processes and institutions that will be flexible enough to respond to new threats, fluctuations in public confidence and changing attitudes and mores. It will not recommend new legal or technological instruments, for reasons detailed in this section.

A feature of this area is the speed with which privacy-threatening technical developments occur, and the relative tardiness of responses to these. As noted above (section 1.2.4), it is unhelpful to view this issue as primarily a legal issue demanding a legal response, and as implied in the same section, neither is this just a technical issue requiring some kind of cleverer widget or protocol to sort things out.

In the legal case, a new instrument (say, for instance, a replacement for the EU's Data Protection Directive 95/46) would simply take too long to craft, given the speed of technical development. One serious problem with 95/46 is that its intellectual background was the world of the 1980s in which the digital threat to privacy was posed by the proliferation of standalone databases within government and industry; its laudable aim was to facilitate ethical data-sharing across the national borders of EU member states to extract value from those databases in the internal market without compromising privacy. However, many commentators have argued that it is completely inadequate for a networked world. Indeed, the problems of the EU lawmaking process in the context of the regulation of technology are rather nakedly illustrated by 95/46, which came into force at pretty well the exact moment that the World Wide Web, whose existence is not mentioned in the directive, became an important social and economic tool beyond the purely academic sector. The problem with legal instruments is that by the time they appear, the threat they are intended to encounter is likely to have evolved.

I note once more that the Directive is currently under revision, having been supplemented already by Directive 2002/58/EC on privacy and electronic communications. The Commission has attempted to move quickly to produce a timely revision, but – illustrating the difficulties here – pressure on the Commission from the French Commission Nationale de l'Informatique et des Libertés (CNIL) and other bodies has already forced a delay in the schedule (Williams 2010).

Technical responses (i.e. software, protocols and tools to protect privacy) tend to become available more quickly (and their developers have a greater technical understanding of the threats), but are problematic in their own right. In the first place, there is still a time lag between the emergence of the threat and the development of tools to counter it. Secondly, privacy-enhancing technologies tend to be hard to use (Sasse & Flechais 2005), often relying on the individual to deploy an unrealistic degree of understanding of the issues (e.g. requiring the individual to state his or her privacy preferences precisely in some technical language). Thirdly, technologists' models of behaviour are often wildly inaccurate, failing to factor in mistakes, short cuts, ingenuity, laziness, creativity and lack of engagement. The degree of vigilance that purely technical solutions demand often places an unrealistically large overhead of responsibility on the individual. Fourthly, technological fixes or patches tend to

deal with more specific types of attack than do the legal solutions. Consequently, patches can introduce further vulnerabilities, and in any case it is part of the natural threat/response cycle that hackers will immediately begin work to undermine the new solutions.

1.4 Structure of this report

To investigate this question, the report will be divided into three major sections. In the next section, I shall discuss the theory of transparency. In section 3 I shall discuss privacy in the context of transparency, focusing particularly on how privacy and transparency can be complementary rather than antagonistic, and considering how trust in the transparency programme can be created and maintained. These two sections will provide essential background to the specific topic of the technical question of how privacy can be compromised by applying computing techniques to digital datasets (section 4), and how traditional models of data management can no longer be accepted uncritically. Section 5, the final substantive section of the report, makes recommendations about how to address the difficulties that have been outlined.

2 Transparency

This section will set out the theory of transparency, beginning with a description of the philosophy behind it (section 2.1), and moving on to consider the nature of any potential threat to privacy from either theory or practice (section 2.2)

2.1 *The theory of transparency*

The appeal of transparency, which grew out of the success of right-to-know measures and limited, targeted transparency programmes such as that which forced corporations to publish their accounts for investors, can be summarised in the famous quote from Louis Brandeis of the US Supreme Court: “sunlight is the best disinfectant.” In a pleasing symmetry, Brandeis was also a pioneer in the development of privacy rights. In this section I shall set out the philosophy of transparency in more detail.

2.1.1 The philosophy of transparency

The basis of the philosophy of transparency is very simply stated: the government has collected data or information for whatever reason, using the resources and legitimacy it derives from its citizens, and therefore, unless harm could result from the public release of that data, there would seem to be little reason against releasing it to its citizens to make productive use of. One of the potential harms is of course the danger that the privacy of some citizens is compromised.

According to the transparency philosophy, government information and data should be *freely* available. Having been collected using public money, it should, where possible, be open for reuse in order to create further economic value. Government open data should be available to all to avoid monopoly exploitation or rent-seeking by a cabal of data controllers. The barriers to entry should be as low as possible, while publication should be inexpensive and straightforward, with as few bureaucratic overheads and layers of management as possible. Speed and timeliness are important. Regulations demanding bureaucratic oversight will provide an opportunity for the process to be obstructed, or ‘kicked into the long grass’.

This means that a transparency programme should ideally eschew common regulatory mechanisms such as putting the data behind an Application Programming Interface (API), a set of rules that other computer programs must follow in order to access the data. APIs can be used to protect data, for example to operate a regime where only licensed developers or subscribers are allowed access. Such control over the released data could be of value in protecting privacy for reasons that will be explored below, but the processes of control would demand too much bureaucracy, and be off-putting for potential users.

Transparency is underpinned by a mixture of arguments from the left, right and centre of political philosophy, which goes some way to explain its wide appeal. From John Stuart Mill it takes the idea that the serendipitous reuse of data is valuable, because “the utmost possible publicity and liberty of discussion” is an important condition for improving our understanding of the world via critical debate, and “the widest participation in the details of judicial and administrative business” will only be effective if the discussants are well-informed (Mill 1861, chapter 6; see also Mill 1859, part II). From Hayek, it takes the idea that knowledge about an economy or society is distributed across its population, and that therefore individuals are best-placed to judge their own information and data needs (Hayek 1945). And from

egalitarian socialist philosophy it takes the idea that asymmetries of information lead to asymmetries of power, and should therefore be eliminated, a line of thought dating back to Diderot (cf. Diderot [1755]1995).

In this review, I shall take this composite philosophy of transparency as a given. It is certainly defensible and appeals across the political spectrum. I shall attempt to develop ideas for protecting privacy within its context. Were this to prove impossible, given the importance of privacy as a value and a right there would be serious consequences for transparency.

Transparency is now feasible and valuable in a way that it would not have been a few years ago. Not only have public attitudes toward authority altered, but crucially technologies are now in place which make it straightforward to disseminate data widely, and for that data to be reused in new, innovative contexts. Of all the digital technologies that are relevant, the World Wide Web is the most obvious.

Note, however, that the use of the Web subtly alters the justification for transparency. If data are to be released on the Web to the citizens of a nation, then – if the openness criterion is to remain – there is no way to restrict its publication to those citizens. The data become available to everyone, whether or not their taxes paid for data acquisition. Diderot, whose *Encyclopédie* had an impeccable internationalist outlook, would have approved.

2.1.2 Data literacy, representation and intermediaries

To be properly informative, data must be usable. Yet government data, sometimes presented in complex, technical or unfamiliar formats, are difficult to understand, boring to look through and hard to manipulate. Without basic levels of data literacy, how can transparency be empowering?

This question is often raised – see, for example, the flurry of sceptical newspaper comment that accompanied the first monthly release of crime data in February 2011 (e.g. White 2011). Data literacy is indeed an issue, and is generally, like all kinds of literacy, a good thing to be encouraged (cf. e.g. Beetham et al 2009, which talks of digital literacy, or McAuley et al 2011). It would certainly make transparency more likely to empower people.

Yet in a world where data literacy is in somewhat short supply, transparency can still make a difference via the intercession of intermediaries, information entrepreneurs and applications developers who amalgamate data from different sources to present a picture of some state of affairs in real time, possibly via websites or smart phone apps for. Fung and colleagues go further and argue that citizens will underconsume data *unless* such intermediaries re-present it (Fung et al 2007, 121). The role of these intermediaries is to present the data in comprehensible form, and devise services around them. These services may be provided free of charge, but some intermediaries will find methods to monetise them, perhaps via advertising, or a subscription model. It may be that intermediaries are able to contribute, via their services, to the growth of data literacy (which may lead to their being able to access education funding).

An obvious point is that these intermediaries are in a position of some power, because they do not present the data neutrally; they have an editorial slant which they will naturally seek to promote. This is of course true, but while it may be a useful corrective to naïve optimism about a sector of public-spirited hackers, it does not hold water as a serious objection. In the current situation, news about the world is filtered

through a smallish number of media outfits, whose power is proportional to the barriers to entry to the intermediary role. Transparency lowers those barriers to entry, and allows many more entrants into the field of data provision. Honesty and good faith cannot be guaranteed, but transparency at least provides the opportunity for alternative views to be heard.

Digital divides and unequal access to data are also important concerns, but once more it is hard to see how one could address this without open data re-presented by intermediaries except by the draconian and counterproductive strategy of starving the data literate of data. Even if one remained stubbornly pessimistic about the capabilities and intentions of intermediaries, it is surely impossible to argue that transparency makes anything *worse* by increasing competition in the data market.

If there is a lack of data literacy in a society, it does not invalidate a transparency programme. Instead it underlines the importance of a competitive set of creative intermediaries.

2.1.3 Government as an intermediary

There is no reason why government should not also position itself as an intermediary in some cases, although this cannot be the only solution to the problem. Even if government does adopt this role, it is essential that the *data* are released, so that other intermediaries can use them as they see fit.

How, and why, might the government set up as an intermediary? As an example, consider the release of crime data. The data are being released on a rolling basis, but the Home Office also set up a site, police.uk, to present them to the public. This site has two entirely laudable functions. The first is to help create a constituency for the crime data, a group of people interested in what they tell us. Rather than release the data and wait, possibly for months, for intermediaries to emerge and present the data imaginatively, the Home Office's sensible strategy was to launch a crime mapping site, with associated publicity, to increase awareness. The second function of police.uk is to provide an interface between the citizen and the police, for example informing people of who their beat officers are, and when meetings between police and the community are scheduled, thereby equipping the citizen to play a more active role in the policing of his or her own neighbourhood.

So much is good practice. Yet there are two potential pitfalls to this approach. The first is that the government, with its extensive resources, could squeeze out other intermediaries, resulting in a less rather than more varied information market.

The second is that intermediaries focus on the *representation* of the data, not the data themselves. Hence the government might find itself, as both a data provider and a data intermediary, looking at the data in incompatible ways. For example, in the UK there has been a perception that it will be problematic to bring data about *crime* (which are currently available) together with data from the *courts*, despite the immense public interest in knowing how crimes were dealt with by the criminal justice system. Of course it would be problematic to do that *on a map updated monthly*, because of various issues to do with the time lags between criminal activity and court proceedings. But that is an issue with the specific representation type – the map. There is no reason why some creative intermediary could not find some other way of amalgamating crime and court data (possibly with other types of data) to produce an imaginative presentation of the unified datasets to the public. The undoubted

difficulties of *representation* can be taken, falsely, to imply difficulties in data *provision*.

Hence government is unlikely to be a satisfactory intermediary beyond the short term, but does have a role in helping boost awareness and data literacy. Independent intermediaries are likely to be much more creative and aware of demand in their use of data. Furthermore, their demand-awareness will also help set the agenda in the selection of datasets to release, and will allow them to work in partnership with government to ensure data quality and reliability.

2.1.4 Two transparency agendas

The transparency philosophy contains two separate and independent agendas which I call the *accountability agenda*, and the *information agenda*.

- Under the accountability agenda, the aim is to move away from traditional models of accountability of public services. Currently, the accountability mechanisms used for public services are internal to government and/or formal (to use the terminology of Gilbert 1959), using resources and processes generated by government itself (chiefly oversight by civil servants, often driven by targets, ultimately grounded in the accountability of ministers to Parliament). The aim is to move toward informal, external oversight (Gilbert 1959) with direct intervention and participation by citizens, interest groups and the media. The accountability agenda therefore requires sufficient relevant data to be provided to allow considered judgement about the performance of public servants, and especially to allow comparisons to be made over time or between agencies.
- Under the information agenda, the aim is to allow the citizen to develop a rich picture of his or her community, empowering him or her by enabling the negotiation and management of community and environment. Under this idea, the government supplies data that it possesses or has collected which could be of value to the citizen, but because the citizen decides what is of value and how, government cannot predict with any accuracy which data that it holds is of most interest.

These two agendas work together in driving change, giving incentives for government agencies and service providers to move from being bureaucratic organisations, to post-bureaucratic organisations. This distinction between bureaucratic and post-bureaucratic has been bandied about frequently in recent British politics, and some commentators have become suspicious of the jargon. Nevertheless, the distinction has been made with due precision; for definitions of the terms, and the characteristics of these types of organisation, see (Kernaghan 2000).

The accountability and information agendas are not equivalent, and may sometimes be in tension. The differences between them are summarised in the following table.

Accountability agenda	Information agenda
Can be driven from the top down. Expertise is helpful for determining which data are valuable for holding government agencies to account for the functions which they are designed	Better driven from the bottom up, as the citizen has a much better idea of which data are of value to him or her, than the government can have in the abstract.

Accountability agenda	Information agenda
to carry out.	
Supply- focused.	Demand- focused.
The data need to be specific only in so far as they allow the important comparisons to be made.	The more specific (less aggregated) the data the better. Citizens need data about their local community, and all things being equal will benefit from seeing those patterns visible at a low level (street level, neighbourhood level, ward level), rather than the patterns that emerge at higher levels (county level, city level).
Assumes a purpose (holding government agencies to account) for the data.	The released data have no purpose as such. I.e. although they may have been collected by government for a purpose, there is no assumption that the data will be reused by the public for this or any other purpose.
Hence it makes sense, in this context, to talk about unintended consequences of a data release (e.g. a government agency might be able to game the system).	The intended consequence is that the citizen be empowered in some way. Hence the only possible unintended consequence would be that the data release disempowered the citizen (possibly by invading privacy).
The data are likely to be useful in isolation (which is not to say they will not have uses when linked to other data).	The data will become much more useful and powerful when linked to other datasets.
It follows from the above that there are fewer imperatives to represent the data in formats that maximise linkability (e.g. RDF).	It follows from the above that there are stronger imperatives to release the data in linkable formats.
Failure of the public to use the data will be disappointing.	Failure of the public to use the data will not be a problem. They have a right to the data, including datasets that remain unused (in the same way that an unused right of way is, and should remain, a right of way).
Based on a right to hold government to account.	Based on a right of access to the data.
Intended to critique government services, and improve service delivery and efficiency.	Intended to increase the citizen's independence of government.
To be effective, there need to be routes and methods in place to allow the citizen to make a difference to the institutions being held to account.	To be effective, services need to be developed around the data.

Tension between the two agendas arises in particular with the *types* of data selected for release, the methods of *selection* and the levels of *aggregation*. For instance, data about crime are most helpful to the accountability agenda when aggregated to a certain degree, so that patterns at the level of a police force's territory are visible, whereas they may be most helpful to the information agenda when disaggregated, so that one can see specific crime hotspots, for instance down to a particular street corner. Release of data in the form appropriate for one agenda may make it hard to infer the data in alternative forms, and hence there is a tension.

In most cases, the two agendas co-exist perfectly peacefully. For instance, again using the example of crime, it may be that accountability is best served by releases of data about serious crime, while the information agenda provides a stronger demand for data about anti-social behaviour. There is no tension here, because of course both types of data can be released in parallel.

2.1.5 What transparency is not

In recent years, there have been incidents of the unscheduled releases of data which have provoked substantial public debate. It is important to make clear how these incidents differ from the transparency programme – and therefore are beyond the scope of this review. Such events do sometimes have privacy impacts, but these will not be remedied or otherwise addressed by the recommendations below. Examples include:

1. In 2007, two CDs belonging to Her Majesty's Revenue and Customs with only weak password protection, containing sensitive details of child benefit claimants in the United Kingdom, went missing from the internal mail service. Around this time, a series of unintended releases of data, through carelessness and poor organisational management of information, occurred in the British government and in the private sector. Almost certainly, similar incidents continue to happen even though the media spotlight has been trained elsewhere.
2. In 2009, in response to a Freedom of Information request, the House of Commons authorities prepared to release redacted records pertaining to the expenses claims of MPs over a period of years. Before the records were released, however, the *Daily Telegraph* published unredacted copies of the same records which received massive publicity and created an unprecedented scandal.
3. In 2010, the document archive website Wikileaks published confidential documents relating to the conflicts in Iraq and Afghanistan, and detailed correspondence between the US State Department and various of its diplomatic missions.

None of these cases is an instance of transparency. In the first case the data were not released openly or published. If the data are still available to anyone, they will be available only to a small number of people, who either found or stole the original CDs, or who have purchased their contents from the possessor.

In the first and third cases, the data release was not planned, and was partial. The total set of potential data users (i.e. the general public) was not alerted to the forthcoming release, which was therefore not a transparent process. Furthermore, there was no guarantee that the data released was a comprehensive collection of the relevant data.

The data may have contained gaps, and therefore may present a slanted or biased view.

In the second case, the *Telegraph* information was not released according to a well-understood schedule and methodology. The information released in accordance with the original FoI request was so released, but the leaks from the *Telegraph* were intended to drive a news agenda, were released in staggered fashion, were accompanied with sensational commentary and also disregarded the data subjects' (MPs') prior understanding of how the process should work. That does not entail that the official, much-delayed release of the data under FoI was a perfect method of publication, and that the redactions were justified; the argument is only that the *Telegraph's* leak was not in itself part of a meaningful, agreed, accountable and legitimate transparency process.

In none of the cases was the release of information itself a transparent process. The distribution of data was restricted in the first case, selection of information was random in two of the cases, while in the case of MPs' expenses there was no due process.

2.1.6 The forms of resistance

The British government has traditionally been somewhat retentive of data and information. Tony Blair, whose government introduced the Freedom of Information Act 2000, famously berated himself in his memoirs as a “naïve, foolish, irresponsible nincompoop” and “quake[d] at the imbecility of it.” The Minister for the Cabinet Office has argued that transparency will require a “radical culture change for the public sector” (Maude 2010).

Parts of government which do oppose the transparency programme will be apt to use not only direct arguments to stop the process, but also indirect ones about the qualities of the data in order to slow down releases. There are a number of such arguments that can be marshalled here, including the need for accuracy (so that the public is not misled by inaccurate data), and the supposed costs of the process. Privacy is one of these useful arguments; raising privacy concerns can prevent data releases even if the risks are small or easily mitigated by aggregation or other means. Hence an *apparent* concern for privacy on the part of a data provider may or may not reflect a *genuine* underlying concern. The Information Commissioner's Office (ICO) urges organisations not to hide behind the Data Protection Act unnecessarily when dealing with requests from members of the public (cf. http://www.ico.gov.uk/news/current_topics/duck_out.aspx). Though this certainly does not entitle those implementing a transparency programme to brush privacy concerns away (cf. ICO 2011), it complicates the issue in a retentive culture where transparency is alien.

2.2 Is there a threat to privacy from transparency?

How might transparency affect privacy? In this section I shall briefly look at this question, at first from the viewpoint of the philosophy of transparency, and secondly from the perspective of the actual implementation of the UK government.

2.2.1 The *prima facie* threat to privacy

There is of course a *prima facie* issue with respect to privacy in any transparency programme. A transparent government must release data, by definition. What if the

data concern me? If data, of which I am the subject, are released onto the Web, where is my privacy then?

It is worth noting first of all that the privacy which is supposedly at risk from transparency is a somewhat odd brand. The data about which people are concerned are held by government, and are not in possession of the data subject. The subject has little control over how the use of the data. To feel private, the subject must trust the government to prevent access by others, follow best practice in both data protection and data acquisition. As Raab (2005, 285) pointed out, "it is no comfort to a privacy-aware individual to be told that inaccurate, outdated, excessive and irrelevant data about her are encrypted and stored behind hacker-proof firewalls." Privacy here is akin to trust in government.

Nevertheless, there is an obvious point that transparency does impact the individual at the point where he or she interacts with the government. If citizens are to judge whether, say, a particular school is underperforming, then they will need to know the examination results of the children in that school; to judge their police force they will need data about crime and from the courts. Of course these can be anonymised or aggregated into statistics, but the main point is that knowing about government necessarily involves knowing about society.

Even if the data are aggregated, so that individuals do not obviously appear in them, there may be worries that a picture of one's life could be reassembled by someone with sufficient information, resources and patience. As we shall see in section 4, this suspicion is not unfounded. Indeed, even when one does not appear in a dataset at all, one can still have aspects of one's life exposed in an unwelcome way.

As an example of this somewhat counterintuitive notion, consider an extraordinary case in Germany, a nation which is generally more conscious of privacy than Britain. After some controversy, Google Street View, a service which provides panoramic views of locations across the world, went live in Germany despite concerns that people's houses would be clearly visible. After some pressure from the German government, Google agreed to a compromise, that people would be allowed to 'opt out' of the service by contacting Google, upon which images of their property would be blurred out.

So far so good. However, in Essen, transparency activists have been attacking and vandalising the homes of those who have opted out, throwing eggs and writing graffiti ('Google is cool') on the blurred-out houses. This behaviour is not endorsed in any way by Google (BBC Online 2010).

The interesting point is that the compromise was meant to be privacy-preserving, but in this case the activists wished to identify only the (whereabouts of the) houses of those who opposed the aims of Google Street View – an identification which could be achieved from the unblurred photographs of adjacent houses which were handily available from Google Street View itself.

Hence we have an example where *not* appearing in a dataset identifies a person via analysis of that dataset. It is extremely hard to preserve privacy when one is a minority presence in a transparent world. Given that difficulty, we need to understand how transparency affects privacy, where the threats are, and how they can be curtailed in order that privacy retains sufficient protection to allow us to live our public and private lives unhindered by unwelcome interference.

2.2.2 Potential risks to privacy

The theoretical possibility of a threat to privacy by transparency has been discussed. What about the practical prospect in the UK in 2011?

Encouragingly, the Public Data Principles are prefaced with the following clear statement.

“Public Data” is the objective, factual, non-personal data on which public services run and are assessed, and on which policy decisions are based, or which is collected or generated in the course of public service delivery. (Transparency Board 2010)

The Minister for the Cabinet Office has explicitly emphasised this definition of public data as specifically non-personal.

However, despite my profound belief that transparency is transformative in its effects, it is crucial to be clear about the type of transparency we mean. We want to be transparent about anonymised corporate data, we do not want to publish the public's personal details. Personal privacy is the bedrock of a free society, and is as important a guiding principle for this government as transparency. That is why our commitment to making government more open goes hand in hand with our commitment to strengthen civil liberties. So as we move forward with our plans for government transparency, it is imperative that we continue to strike the right balance between this openness agenda and individual privacy. (Maude 2010)

Nevertheless, where the citizen and the state interact, and the performance of government agencies depends on the behaviour of individuals, data about government may indeed include data about individuals. The crime data released on 1st February, 2011 are an example of this kind of case, as are some of the datasets mentioned in the Prime Minister's letter to Cabinet Ministers of 7th July, 2011 (Cameron 2011), particularly with regard to health, education and criminal justice. Definitions of 'personal data' are not sufficiently precise to be of a great deal of help in borderline cases such as these. It is not always easy to tell whether someone is identifiable from a dataset, or from a combination of that datasets with others (some of which might be very rich sources of information, such as newspaper reports, or material placed on social networking sites). The grey areas here are very grey; it is not a simple matter to examine data and determine whether or not data subjects are identifiable.

We should also note that the ultimate aim of the transparency programme in the UK is to move from the current top-down regime to a more demand-driven one where citizens can express their interest in particular datasets. Because the initial programme was rather more focused on the accountability agenda than on the information agenda, there was a corresponding focus on comparable data that express something about how well a department or agency is functioning. Personal data could easily be selected out. If we move to something closer to the information agenda, then citizens will begin to demand data that is meaningful for them – which will naturally include some data that could be personal, including health data, educational data and criminal justice data (Cameron 2011 has moved in this direction). Of course such data will not and should not be released *automatically* and will be anonymised, but there will be debates in the foreseeable future in which privacy and transparency will be in immediate tension. In such cases hard decisions will need to be taken about whether

the public goods which are hoped to follow from data releases can be sufficient to justify invasions of privacy.

In his latest letter (Cameron 2011), the Prime Minister committed to release a number of anonymised datasets in open standardised formats under the Open Government Licence, many of which are tailored to the information agenda and which will be very valuable to communities. Of these commitments, perhaps the most striking from a privacy point of view is “opening up access to anonymised data from the National Pupil Database to help parents and pupils to monitor the performance of their schools in depth, from June 2012. This will enable better comparisons of school performance and we will look to strengthen datasets in due course.”

I do not wish to go into detail about the datasets considered for release in (Cameron 2011); as I argue elsewhere, whether a particular data release is unacceptably risky should be determined by relevant experts in context, not by crafting general-purpose rules in advance. However, there is no doubt that the 2011 commitments will require the release of anonymised datasets, and there are issues concerning the effectiveness of anonymisation and the levels of risk of deanonymisation, which will be explored further in section 4 below.

3 Trust, privacy and transparency

Any transparency programme needs to be sustainable. For this to happen, it will need to retain the confidence of the citizenry. In any system where data are required to be provided on a regular basis, and where data consumers interact in a relatively sporadic and infrequent way with the system, there is a danger that over time the system will start to serve the providers' interests rather than the consumers' (Wilson 1980, Fung et al 2007, 11, 106, 109). Citizens need to know that transparency serves them as well as or better than it does the providers, in order to have sufficient confidence in the system to make it a success. They also need to be confident that the data provided are as complete and accurate as possible, and to be sure that the data are more likely to be used for their benefit than to their detriment.

There are many important factors in the maintenance of trust, but privacy is central to the concerns of this review. If citizens come to believe that an effect of the release of public data will be a significant decrease in privacy, then the result will inevitably be a withdrawal of support and a reduction in the democratic legitimacy of the programme.

It follows that the maintenance of privacy, as the Minister for the Cabinet Office has argued, is essential to preserve confidence, and the best way to do that is to embed privacy protection within the programme itself, rather than having it as a bolt-on component. If the processes underlying the transparency programme are clearly privacy-preserving (and seen to be so), then public confidence (in that aspect of the programme at least) should be retained. Far from being in opposition to transparency, privacy is *necessary* for it.

In this section, I shall explore this issue in more detail. I shall begin with a discussion of public attitudes to privacy, and then consider the role of political and technical debate about privacy in fostering confidence in transparency. Legal redress for breaches of privacy is of great importance, and will be discussed in section 3.3. Section 3.4 discusses the role of transparency about transparency in creating trust and confidence. Finally, section 3.5 raises the question of whether transparency might be aided without damage to privacy if there were different classes of personal data.

3.1 The citizen's reasonable expectations

3.1.1 Privacy and context

Public attitudes to privacy can appear confused or inconsistent. Surveys show a great interest in privacy (e.g. Bradwell 2010), but actual behaviour (for example, using online social networking sites) often reveals a rather cavalier attitude, as people give away and publish intimate details of their lives (e.g. Joinson & Payne 2007). These apparent confusions, however, can often be explained away as caused by shifting contexts, ignorance, inappropriate levels of abstraction in presenting a problem to people, and varying attitudes to privacy, in particular between generations. Let us briefly consider these in turn.

Context. Privacy is not only about what information is around, but many other matters connected with the social situation (Nissenbaum 2010, and see section 3.1.4 for more detail on her work). Of course we have varying attitudes about who knows what – I might be comfortable with my doctor knowing fact A, my bank manager knowing fact B, and my friends knowing fact C, but extremely uncomfortable with

anyone knowing the conjunction of these facts. But context goes further than this – even the method of transmission of information is important. If I ask my wife what she is doing tomorrow, she will tell me (I complacently assume). She is comfortable with my having that information. However, if I didn't ask her, and discovered the same information by looking in her diary without permission, she might well be justifiably annoyed. And if I sent a FoI request to her PA for the information, relations might become extremely tense, not to say glacial, at home. Yet it is the same information I receive each time; the relevant difference is the method of acquisition.

Ignorance. In 2010 the websites *pleaserobme.com* and *icanstalku.com* achieved notoriety by raising awareness that people were inadvertently sharing information about their whereabouts using Twitter and releasing geotagged information (for instance photographs with the location embedded in the metadata). Someone who tweeted that they were at a lecture or a concert, thereby revealed to anyone following their Twitter feed that they were not at home. Someone who placed a photograph of their own home on a photo-sharing site such as Flickr, and tagged it as such, could reveal where they lived (since the location of the camera, sometimes contained in its metadata, would of course be close to the location of their home).

In other types of case, people would place revealing photographs of themselves and their friends online, and describe aspects of their lives on social networking sites such as Facebook. Surveys (e.g. Karyda & Kokolakis 2008) have shown that people sometimes had very little awareness of the implications of their actions – for instance, it is not widely appreciated that Facebook owns the data on its site, and that its most promising business models are based around the targeted advertising that access to those data makes possible.

Abstraction. Any breach of privacy takes place in a specific context where there are often benefits that can be achieved by sacrificing privacy. For instance, a mobile phone gives one's position away as it locates and interacts with the nearest antenna; this is naturally privacy-invasive as it tells the phone network where one is (for example, it has recently been revealed that Apple iPhones and 3G iPads have been storing data about the devices' whereabouts in unencrypted form – Allan & Warden 2011), but the convenience of mobile phones means that people are prepared to discount that worry. People collect points with store cards, even though it means that supermarkets are able to build up detailed databases of their likes, dislikes and patterns of purchase. An intimate blog may reveal many details about one's personal life, but the gain in self-expression and assertion is usually felt to compensate (interestingly, many explicit bloggers are embarrassed when their friends or family read their writings, but have no problem with total strangers – de Laat 2008). In the abstract, the issues surrounding location-based privacy, or commerce, or intimate blogging, seem remote and hard to comprehend. In some (perhaps only a few) particular circumstances, however, one's behaviour may lead one into unanticipated, though avoidable, problems.

Attitudes. It is also widely thought that there is a generation gap in attitudes to privacy. Some believe that younger people have no interest in privacy, but this is false. Rather, there is a difference in interpretation of privacy.

- Among older people, a *control* model predominates (and this is reflected in most privacy law). The person feels private when he or she is able to control others' access to information about him- or herself.

- Among younger people, an *anonymity* model predominates. The person feels private if freely available information cannot be connected with him or her (cf boyd 2008).

Younger people (digital natives, to use a common term) have become very experienced in negotiating the benefits and pitfalls of sharing information for social ends. It may be that over-enthusiastic information sharing is unwise and will tend to get out of one's control (Bailey & Kerr 2007), but even people who have revealed major aspects of their lives online usually have an interest in preserving their privacy to an extent that suits them.

In summary, privacy makes sense only in a particular context. Without knowing what information may or may not be shared, with whom, for what reason, in what form, and for what potential benefit, privacy is a meaningless abstraction – something, like motherhood and apple pie, that virtually everyone is in favour of.

3.1.2 Use Limitation

It follows from this that one way of making decisions about privacy is to fix the context of data. This principle is enshrined in a number of privacy codes, and is often referred to as the use limitation principle. For instance, an influential set of OECD guidelines (non-binding, but which have been incorporated into a number of binding statutes and conventions over the years) published as long ago as 1980 (OECD 1980) set out the Purpose Specification Principle and the Use Limitation Principle.

Purpose Specification Principle

9. The purposes for which personal data are collected should be specified not later than at the time of data collection and the subsequent use limited to the fulfilment of those purposes or such others as are not incompatible with those purposes and as are specified on each occasion of change of purpose.

Use Limitation Principle

10. Personal data should not be disclosed, made available or otherwise used for purposes other than those specified in accordance with Paragraph 9 except:

- a) with the consent of the data subject; or
- b) by the authority of law.

There are privacy-preserving merits about such principles, as they provide the data subject with the context in which data will be used. If personal data are used other than for the original purposes, then either the new use has the force of law behind it (in which case the subject would in any case have no right to prevent it happening), or will have to be explained to the subject when his or her informed consent is sought. It has been argued that use limitation is increasingly important in the protection of privacy (Brown 2010).

However, the application of these principles is not as simple as one might hope.

3.1.3 Transparency and context

It is obvious that the principles do not sit very well with the ideology of transparency. The whole point of transparency is that serendipitous reuse of data is powerful – one cannot predict all the circumstances in which data will be valuable. Hence, when data are published on the Web, not only can one *not* say how they will be used, the *whole*

aim is for them to be used in an unexpected way. Against the Use Limitation Principle, the transparency ideology places a Use Maximisation Principle for non-personal data (recall, for instance, in the UK public data are specifically defined as non-personal data). Transparency demands that the contexts in which data are released are unlimited by terms and conditions, representational and formatting issues, and so on. Because the Use Maximisation Principle of transparency applies to a different set of data to the Use Limitation Principle, there should be no clash. If data are personal data, then the Use Limitation Principle will trump use maximisation.

However, there will be grey areas. The transparency argument that there is a great public interest in the unlimited dissemination of data should not be dismissed out of hand, especially if it is unclear whether a dataset is personal data or not.

In short, the philosophy of transparency actively seeks out new contexts for information and data. To the extent that such a philosophy is accepted, instruments such as the Use Limitation Principle will be problematic; to the extent that purpose specification and use limitation are adhered to, transparency is handicapped.

3.1.4 Context and expectations: the theory of contextual integrity

Failure to understand the importance of context to privacy, and the potency of context-relative expectations for privacy perceptions, is one of the main explanations of the difficulties in developing privacy policies for new technologies and institutions, according to academic commentator Helen Nissenbaum (2010). Nissenbaum's theory of *contextual integrity* is a useful tool in this space. This recommends specifying the social context in which a data management practice takes place, understanding the norms, expectations and actors which are active in the context, and specifying the data and information flows characteristic of a new practice – in other words, what do people expect to happen, and what actually will happen when a new information management practice (in our case, the release of datasets onto the Web with little or no restriction) is put in place?

If we think of transparency using this framework (cf. Nissenbaum 2010, 182-183), we find some constant factors. New methods of transmission of data by publishing datasets online for the world to download bring a whole new set of actors into the context – all those people who were previously unable to get the information because it was held by government. Other aspects of the context are broadly unaffected.

Two sorts of evaluation are required. The first must issue from the context itself: will the data release undermine the aims and goals of legitimate actors in the context? To take the example of the crime data released monthly from 2011, will the release of crime data make the police's job harder? For example, will people, knowing that 'their' crime, of which they were a victim, will appear on a crime map, be reluctant to report crime for some reason as a result, thereby causing degradation in the quality of crime data themselves? If so, then a rethink will be necessary.

The second is a wider examination of the moral and political factors in operation in the particular context. In particular, could the data release cause a reduction in autonomy or privacy in this or other contexts? Could it cause unjust outcomes? Could some agent gain unwarranted power via processing of the data? Will people's reasonable expectations of privacy in that context be completely undermined, so that they are acting under a set of expectations (e.g. confidentiality) that are now false?

Understanding the expectations and norms in any particular domain enables the maintenance of the contextual integrity of the situation. This will be extremely important for retaining the legitimacy of the transparency programme, by ensuring that no radically new data management practices are imposed upon an unwilling and unprepared public. However, gaining that understanding is not as easy as it might sound. At a minimum, those pushing forward with transparency need to be highly aware, and not dismissive, of public opinion, and need to devise process management institutions which include widespread consultations and sample a range of opinion (in the recommendations to this report, I shall recommend a process structure that brings in such a range).

3.1.5 Empowerment and consent

One desirable feature of data management systems is that they should empower the subjects of data where possible. The obvious method here is via consent of the subject about the treatment of personal data about him or her. Consent is an important part of the Use Limitation Principle.

However, consent is something of a blunt instrument. The current model of consent is somewhat unsatisfactory – for example, on a website, one is presented with a link to a long, complex and tedious privacy policy, and one ticks a box to show that one accepts it. Few are capable of understanding such policies, and even fewer bother to read them. Yet by giving consent to one's data being used in this very binary fashion, one may be signing up to more than one bargained for. To take one example, Facebook's privacy policies shifted from an opt-in to making data available to third parties, to an opt-out – yet it was not required to ask for its users' consent for this significant change. The original consent for a totally different policy is deemed adequate.

Furthermore, consent may well not be such a black-and-white issue. There are nuances – one may consent in some circumstances to the use of one's data, and not in others.

There are good reasons to explore making consent a much more useful tool in the context of transparency. At the moment, the binary yes/no model is not very useful. A more nuanced approach, that would not only allow a subject to differentiate between contexts, but also allow the revocation or reinstatement of consent even after a dataset is released is an attractive option, if at all possible (Whitley 2009, Bonnici & Coles-Kemp 2010, Coles-Kemp & Kani-Zabihi 2010).

At the moment, however, the bureaucratic overhead of managing such a system would be too great for a transparency programme to undertake. As noted in section 2.1.1, the management of transparency has to be as lean as possible. Nevertheless, there are research projects in the UK currently looking at mechanisms, institutions, formalisms and architectures for supporting detailed and sophisticated consent management (cf. Encore, ENSuring COnsent & REvocation, <http://www.encore-project.info/>, and VOME, Visualisation and Other Methods of Expression, <http://www.vome.org.uk/>), and the outputs of these projects should be considered by future managers of departments' transparency programmes. Currently, this research is not at a mature stage, but it is an important and promising direction in which to go.

Hence, as current consent mechanisms are (a) not very empowering, (b) bureaucratic in operation, and therefore (c) unsatisfactory, I shall not include any universally-binding recommendations about including consent in this review (although it may be

appropriate in some contexts which I will try to specify – see recommendation 1). However, this is a promising direction for research, and the state-of-the-art should be kept under review.

3.2 The role of debate

Given the somewhat difficult state of the law with regard to data protection (which as we will see in section 4 will get even murkier, once we see how hard it is to decide what is and what is not 'identifiable'), there is an enhanced role for debate and discussion to play in the transparency programme. Complex legal and technical argument needs to be amalgamated into a cogent decision-making process, which will be important evidence about whether government has been properly careful of, or alternatively negligent with, citizens' privacy.

3.2.1 An ideal structure

Ideally, a transparency programme would be driven by information providers with strong interests and incentives to publish. A government department would wish to publish some data, perhaps in response to requests from the public. It would examine the case internally (and transparently), and be enthused about the public interest in putting the data out. Meanwhile, defenders of privacy would be invited to scrutinise the proposed release. If there was a *prima facie* case for redaction, aggregation, further anonymisation or even scrapping the release altogether, the department and the privacy activists should debate it out in some democratically legitimate forum. In short, information providers would be transparency enthusiasts, and they should be made to prove their case under rigorous scrutiny.

This, sadly, is unrealistic in the general case, though a model exists in the Office for National Statistics' Microdata Release Panel (http://www.statistics.gov.uk/about/ns_ons/ons_microdata_releases.asp), which scrutinises requests for access to unpublished microdata. It should be noted that these releases are not transparent in the sense used in this report. In order to protect the confidentiality of the data subject, releases are governed by confidentiality agreements, are available only to recognised institutions and appropriate individuals and for specific statistical purposes, so that the ONS retains a measure of control.

As noted, privacy (alongside other convenient arguments against transparency, such as national security and accuracy) is likely to be used as a delaying tactic. A government department which was less than enthusiastic about transparency would be less than committed to assembling the best case for a data release. The history of transparency shows that only rarely do information providers become enthusiasts (cf. Fung et al 2007), and that they are as likely to try to subvert or manipulate the system as to take part freely and openly. One naturally anticipates a culture change in Whitehall, but there is no reason to think that it will be either speedy or consistently spread across departments.

Nevertheless, the key notions of the envisaged debate – rigour, transparency, detailed scrutiny and advocacy for both transparency and privacy – are important and attractive in this context.

3.2.2 Content of the debate

What would the debate need to cover? As we shall see in section 4, there can be no guarantees about the safety of a particular data release, even of anonymised or

aggregated data, and it is undesirable to continue to conduct our affairs in this space under the pretence that it is.

The issues that therefore need to be addressed in any decision about a data release onto the Web are the risks of release, and the potential benefits. The debate, therefore, should be a risk/benefit analysis.

It is important that the analysis is not 'captured' or driven by a particular interested party, and that the outcome is not pre-determined. Such a debate therefore demands the representation of diverse stakeholders, ideally including those demanding the data, and who therefore can provide some context for the discussion. In recommendation 7, I shall suggest a list of the stakeholders who would ideally be present.

Nissenbaum (2010, 181-183) provides a decision heuristic, based on her theory of contextual integrity, for deciding on whether to introduce a technological innovation, and if so how to regulate it. Those taking part in the debate may or may not subscribe to her theory, but her heuristic provides a good structure independently of whether her theory is being followed. She advocates looking at structural aspects of the context, and then at the following issues:

- A *prima facie* assessment.
- Evaluation of the moral and political factors affected by the data release.
- Evaluation of how the data release affects the goals and values of the context.

The debate could either result in a decision, or a report. Whoever did make the decision would need to compare risks and benefits and decide accordingly. What seems correct (and acceptable to the public) will vary depending on social, political and intellectual attitudes of the day. I therefore do not wish to suggest particular rules, restrictions or constraints at this point.

3.2.3 Debate and audit

A detailed debate between those proposing a data release and those opposing has much to commend it. In the first place, open debate would enable the relevant points to be put and tested, and would be more likely than any other forum to uncover and quantify potential risks to privacy, and benefits of transparency. Debate will enable a more accurate and objective model of threats and opportunities to be constructed.

Second, transparent public debate (e.g. with detailed minutes published online) will facilitate media scrutiny.

Third, visible evidence of careful attention to detail and protecting privacy would help the public appreciate the commitment within the transparency programme to the protection of its interests. Confidence in the programme would be more likely to be maintained if it were clear that issues pertaining to privacy were taken seriously, and embedded in the process itself.

Fourth, in the event that a data release came before a court, the rehearsal of the debate would clearly be valuable for the government in constructing its case. The involvement of experts in debate and scrutiny would mean that the important arguments would have received an airing. In the event that a data release was being contemplated that could possibly be challenged under Article 8 (which is unlikely in the near future, but which cannot be ruled out absolutely), the debate would help determine which of the Article 8 qualifications (national security, public safety,

economic well-being etc.) was being invoked. The very existence of a scrutiny process that was prior to, and as rigorous as, the legal process would anyway help lower the risk of legal proceedings being brought.

Fifth, whereas a legal debate is focused very much on legality, a debate that brings in experts from beyond the legal discipline together with lawyers will be of much greater value in promoting confidence. The use of technical experts who understand procedures for deanonymising data, for instance, would inform any discussion of the likelihood of data subjects being reidentified. As noted above, it is important for the success of transparency that it retains the confidence and trust of citizens; if the government's actions are morally dubious, technologically uninformed, or against the contextual norms and expectations of citizens, the fact that they are nonetheless legal will hardly promote trust.

Sixth, such a debate will help promote awareness and disseminate best practice among privacy professionals. This would be of particular value for Chief Privacy Officers and managers of small to medium-sized enterprises who have to deal, on slim resources and often with little training available, with difficult issues about data sharing. It may also help promote consistency across the sector.

The account of the debate and its outcome would resemble a forensic Privacy Impact Assessment in many ways (PIA – for the ICO's advice on PIAs, see http://www.ico.gov.uk/upload/documents/pia_handbook_html_v2/index.html). Its publication would be an important resource for those with interests in the application and regulation of transparency and privacy.

3.3 Harms and remedies

The approach of European privacy law, attempting to anticipate and regulate harms in advance, is problematic in this sphere, as the nature and extent of those harms are so unclear. In the US, the existence of specific privacy torts (Prosser 1960, Soma & Rynerson 2008, 31-43) has the advantage of allowing releases of data under a transparency programme, but the disadvantage of doing nothing to address risk (at least until a court has ruled on a specific existing problem). There is a difference of opinion about the existence of a privacy tort in the UK; Wacks (2006) argues that judges are not interested in privacy, and tend to use other doctrines such as breach of confidence as a poor substitute for a privacy tort, while Phillipson (2006) maintains that actually breach of confidence has morphed into a privacy tort properly so-called, although even this falls far short of the protection that the European Court of Human Rights seems to assume. Both, however, seem to concur that privacy is strongly associated with breach of confidence in UK law, and that the protection is not very great.

Certainly any transparency programme worth the name needs to consider the possibility of harms resulting from the release of government data onto the Web, and means of remedy.

Insisting on citizens' having recourse to the courts seems entirely unsatisfactory. This is generally a time-consuming and expensive process that would be extremely difficult to undergo for an ordinary person (even celebrities have been known to struggle). Even with legal aid or no-win-no-fee agreements, this would be a tough and off-putting course of action to have to take.

Furthermore, the release of large datasets raises the possibility that many cases may hit the courts at the same time if the privacy of a number of people is compromised.

Finally, the possibility of jigsaw identification of a data subject means that it may be very hard to prove that a government dataset breached a citizen's privacy rights, as the identification might have used data from a number of more or less innocuous sets. The government data might not have been especially worrying – as we shall see in section 4, the science of deanonymisation has made impressive advances, and information valuable to an adversary can be constructed from apparently unproblematic data.

Currently, one role of the Information Commissioner (IC) is to administer fines for breaching the Data Protection Act. It can also investigate complaints from the public, although it cannot award compensation. The IC should therefore remain the judge of whether a data release under the transparency programme incurred too great a risk – in making this judgement, the debate trail discussed in section 3.2.3 above should be a valuable tool for the IC.

One function that the IC currently does *not* have is the ability to compensate those who have had privacy breached. This may not be too important an omission, as a privacy breach may not lead to direct financial loss for a victim – rather the harm would be less tangible and arguably irremediable by financial means.

The needs of a future where transparency was an assumption of government are hard to judge in the absence of evidence about what such a world would look like. Hence it would seem that there are two questions which will need to be kept under review.

- How many cases are reaching the ICO? It may be that the ICO is swamped by cases as a result of data releases about a substantial number of people. In that event, there may need to be new rules, or even a new institution, to prevent the ICO being overwhelmed. One has also to note that if many aggrieved citizens were moved to complain, it would be hard to see how public confidence and trust in transparency could survive.
- Are there significant levels of financial loss or other types of loss as a result of privacy breaches for which a just remedy would be financial compensation?

3.4 Metatransparency

I have made the argument that the UK's transparency programme depends on retaining the trust and confidence of the public. This depends on ensuring that citizens' interests, including but not restricted to privacy interests, are central to decision-making. It also follows from this that citizens need to be sure that transparency is not a way of pushing one particular political agenda, and that the data released under the transparency programme are chosen for their value to citizens both for holding government to account, and as a means of gaining a rich picture of their community.

To this end, the most obvious principle that should attach to transparency is what we might call *metatransparency*, or more prosaically transparency about transparency. The decision-making about transparency should be as clear and open as about any other government business. This entails a number of things, of which perhaps the most important are:

- Transparency of *data assets*. Releases of data are generally welcome, and the data should be accompanied by metadata contextualising and explaining the

principles underlying the data. It is also important for citizens to know what other relevant data have been collected, whether those data are available, and if not why not. In this way, consumers of data will be able to make a judgement about whether the data they have got give a clear or biased, full or partial picture. Datasets should be described, but need not be released if there is good reason not to (e.g. to protect privacy).

- Transparency of *debate*. As noted in section 3.2.3, there will need to be discussions about the dangers to privacy from the release of some datasets. These debates, if conducted in good faith and released in full, will be very important for increasing trust that citizens' privacy is being properly protected, and that data releases are being properly scrutinised.

3.5 Personal, public and private

Nissenbaum's discussion in terms of context helps explain why it has proved extremely difficult to draw a neat line around private space, to delineate it from the public, and why when the situation changes (as for example when new technologies appear), previously unexpected fuzziness or underspecified boundaries are revealed (Nissenbaum 2010, 89-102). It is very difficult, in the absence of a specified context, to determine exactly what is public and what private; there are many clear-cut cases, but the boundaries are very hard to draw. It may be that the publication of certain types of personal data would be less harmful, and perhaps (given a public interest argument) even welcomed by the public.

This sort of distinction is made by the German courts, which recognise three spheres of personality, the *intimate*, the *private* and the *individual*. The intimate sphere includes thoughts and feelings, sexual behaviour and health information, and receives maximal protection. The private sphere, which includes information about one's family life, gets qualified protection. The individual sphere, which relates to one's public, professional and economic life, gets the lowest level of protection (Wacks 2010, 67-68).

However, there is a strong doubt uncontroversial, non-damaging personal data and more intimate personal data can be treated differently. Recent discoveries in the world of computer science and computer security have effectively ruled out the distinction. These discoveries, and technological responses currently under exploration, will be discussed in the next section.

4 The science of deanonymisation

In this section, I shall examine the use of technology to deanonymise datasets that contain personal data. In recent years, the anonymisation and pseudonymisation of datasets have been used as important methods to preserve (much of) the utility of data, and allow it to be shared (or sold) without incurring bureaucratic costs of compliance with data protection laws. However, this broad strategy is under challenge; various theoretical results have undermined assumptions about the safety of anonymisation and pseudonymisation. I emphasise that these results are theoretical, and that there is very little empirical evidence that this theoretical risk translates into a real-world threat, but it is nevertheless a serious issue for any transparency programme, which can only thrive to the extent that the data that it produces has utility. It could be a serious constraint if utility has to be severely curtailed to preserve anonymity because of the dangers of identification of data subjects.

Alternatives to anonymisation will be considered, but they will also prove problematic in the context of transparency.

4.1 Anonymise-and-forget

Anonymisation has been used for a long time as a means to enable data controllers to share data without revealing the names of data subjects, whether within their own organisations or with the outside world, and whether done *pro bono* or as a commercial transaction. With regard to transparency, it is clear that anonymisation, if it lived up to its name and genuinely rendered the data subjects anonymous, would be a powerful tool for a transparency programme, which could release valuable and desirable datasets about people with little or no risk of their reidentification.

4.1.1 Anonymisation

To anonymise a dataset is to manipulate the data in order to make it very difficult to identify the data subjects. Typically this is done by diluting the information content, making the dataset less specific. The most obvious ways to do this are to remove identifying information from the set (suppression), to aggregate it, or and to abstract.

So, for example, if we had an entry in a database that read:

Name	Age	Salary	Profession
Ichabod Thrushaven	37	£30,000	Logistics services manager

Then an example of suppression would be to delete the cell containing the subject's name (or to replace it with some non-connoting identifier, such as an arbitrarily assigned reference number). Examples of abstraction would be to express the subject's age as part of a range (e.g. 31-40), or his profession as a predefined category, such as 'managerial'. Aggregation means to take the data (e.g. the salary data), and release it in some kind of aggregate form, such as the average salary of all the people in the dataset, which is not associated with any individual subject. Of course all these types of change reduce the value of the data in the dataset, but they help preserve the privacy of the subject.

A variant of anonymisation is *pseudonymisation*, where the identifiers in the data are replaced consistently with artificially-generated identifiers. This allows connected data entries to be linked together, preserving the value of the connections. For

instance, in the above example, the name 'Ichabod Thrusthaven' could be replaced throughout the data, wherever it appeared, with the same numerical identifier, randomly-generated especially for the purpose, e.g. '39484'. Then the data would allow certain limited inferences to be drawn – for instance, that the person who attended a clinic on 30th January for alcoholism was *the same person* who saw a consultant on three occasions in July for memory problems – clearly a valuable resource, even if the data do not allow the data processor to identify the person by name.

A further method of protecting privacy is to *perturb* or put noise into the data – deliberately changing some values while keeping important aggregates constant. So, for instance, Ichabod's salary could be given, falsely, as £40,000, while someone else's could be given as £10,000 less than it actually is, thereby keeping the average salary of the subjects of the data the same.

In the rest of this report, I shall concentrate on anonymisation, although the same issues pertain to pseudonymisation (the AOL dataset discussed in section 4.1.3 below was pseudonymised, for example). Furthermore, one can assume that any reasonably sophisticated anonymisation or pseudonymisation method will include perturbation of the data. Broadly speaking, anonymised data are less useful than pseudonymised data, because one cannot tell, for example, whether two separate pieces of data apply to the same data subject. For the same reason, however, anonymisation is less risky. However, with respect to the arguments discussed in this report, the distinction is not important, and I shall use the term 'anonymised' to mean data from which the subject is not intended to be identifiable.

4.1.2 Sharing anonymised data

The idea of anonymisation led to the doctrine/method of *anonymise-and-forget* or *release-and-forget* – a data controller can get maximum *social* value from the data he or she controls by anonymising it and then putting it out into the outside world (whether for money or not), allowing others to extract value from it. Once the data have left his or her control, then it is no longer their problem.

Under this regime, data which may identify data subjects are anonymised, removing information from the dataset so that people cannot be identified from it. Then the data could be released, sold or published for free download without any privacy or data protection issue.

Legally, sufficiently rigorous anonymisation has been held to get around data protection law. If the identifiers are removed from a datum, then it is deemed that the subject is not identifiable, and that therefore it is not personal data. Where anonymisation has not been robust, or when it has not been adequate, there have conversely been legal problems. For instance, Anderson (2008, 294-295) gives an example of health service data in the UK in which people were de-identified and then re-identified later, so that people who had participated in the data collection under assurances of anonymity had been deceived. Other datasets were kept, supposedly anonymised, but with date of birth and postcodes intact, making the subjects highly identifiable.

There are obvious attractions to this model, which has been used for many years, for example in the health sector, without serious compromise of privacy (Kelsey 2009), but sadly, recent work in the field (in particular Ohm 2010) has shown that *anonymise-and-forget* provides less protection from someone trying illicitly to

reidentify data subjects than has previously been thought (I shall follow usual practice in calling such a person an *adversary*). In this section, I shall describe the anonymise-and-forget model, and in section 4.2 discuss its flaws (which, to be fair to data controllers, have only recently become obvious).

4.1.3 Failures of anonymity

Merely anonymising data may not be enough to make subjects unidentifiable from the data in the dataset considered in isolation. In the United States, Latanya Sweeney (1997) proved some years ago that information that may not appear to be identifying can, taken together, be very potent.

I conducted experiments that demonstrated how de-identified health data can be linked to a population register in order to re-identify by name the persons who are the subjects of the health information. Using the voter list for Cambridge, Massachusetts, I showed how a few demographics combine to uniquely identify individuals. It was found that 12% of the 54,805 voters had unique birth dates (month, day and year of birth). Therefore, any information on these individuals that included birth date and city, would almost certainly be specific to the named individuals. Further, birth date and gender together were unique for 29%, birth date and a 5-digit ZIP (postal code) were unique for 69% and birth date and the full 9-digit ZIP were unique for 97% of the voters. These results demonstrate that combinations of characteristics can combine to construct a unique or near-unique identifier which is termed a *quasi-identifier*. These results further show that the typical de-identification technique applied when releasing information for public-use in the United States, does not render the result anonymous. (Sweeney n.d., 21)

Many datasets contain quasi-identifiers as Sweeney discusses here, such as postcodes (possibly abridged, with the first three or four digits), gender and birthdate; these will raise the risks of identification.

As an example that occurred with a pseudonymised dataset, in 2006, the American Internet services company AOL released – onto a website for download by researchers into online behaviour – 20m search query terms typed into their AOL search engines by 650,000 users over three months. This was a valuable resource for research, and AOL was trying to be public spirited by making the release. The data were anonymised, with unique identifiers (randomly chosen numbers) replacing identifiers such as the AOL username, or the IP address of the user's computer. The reason identifiers were provided at all was so that researchers could correlate the different search terms used by the same individual, without their being able to find out who that individual was.

However, search terms themselves are very identifying. Many of us, possibly for reputation management, possibly out of Narcissism, search for ourselves on the Web. If someone searches for "Ichabod Thrushaven" more than a couple of times over a three month period, then, if the person concerned is not particularly famous, it is a fair bet that the searcher is Ichabod Thrushaven himself. His entire search history – all the information he has attempted to find out from the Web – is then laid bare. This may include medical information, information about sexual services, information about political or religious groups and other types of information that is deeply personal. In the AOL case, certain individuals were identified from the data release (Anderson 2008, 295, Ohm 2010).

4.1.4 Data and their use

There have been some responses to this issue – for example, the development of the notion of *k-anonymity*, in which someone is *k*-anonymous in an anonymised dataset if the data about them is indistinguishable from at least *k*-1 others in the dataset (Sweeney 2002). In other words, there is a group of *k* indistinguishable people containing the data subject in the dataset at the minimum. To anonymise a dataset using *k*-anonymity, the data controller needs to delete or suppress the minimum amount of information that will ensure that everyone in the dataset melts into a group of *k* or more people.

However, even *k*-anonymity can be attacked, and is particularly vulnerable when background knowledge can be brought to the deanonymisation effort, and when the sensitive values in a group of people lack diversity. As an example of lack of diversity, suppose the sensitive attributes are health status and sexuality (in other words, we are keen to ensure that information about health and sexuality does not leak out). And suppose that a person is made *k*-anonymous, and is indistinguishable in the anonymised dataset from *k*-1 other people. However, if all or most of the people in that group of *k* have similar health status or sexuality – so there is little diversity in the sensitive values of health and sexuality *across the group* – then it will be easier for the adversary to undo the anonymisation and discover the sensitive information.

This kind of approach assumes that it is possible to make a clear distinction between identifying and non-identifying data, whereas actually the difference is made by the use to which data are put. In an era of scarce data, the distinction between identifying and non-identifying data may have made some *de facto* sense, but now there are very rich sources of information about several aspects of the lives and interests of very many individuals, any data that distinguishes one from one's fellows can be identifying. The list of identifying or quasi-identifying attributes cannot be fixed in advance. Some attributes are identifying in their own right – the name, to take an obvious example, or various biometrics – but *any* attribute, *in combination with others*, can be identifying.

4.2 Jigsaw identification

Hence, many databases do contain the materials for an identification of data subjects, if supplemented by fairly straightforward information from sources such as the electoral roll, or local knowledge. If the adversary has access to rich information resources, then the problem of so-called jigsaw identification escalates dramatically. A number of *coups de théâtre* by academics and privacy activists in recent years has left this beyond doubt.

Note the obvious point that almost everyone in practice does have access to rich information resources via the Web. Transparency will contribute powerful sets of government data to the mix. This section will briefly review the ways in which data can be deanonymised.

4.2.1 Identification with supplementary information: the example of the Netflix Prize

In 2006, the online DVD rental company Netflix released data about 500,000 of its users' movie recommendations. The company sponsored a \$1m prize to be won by anyone who could take the test data, and come up with an automatic movie recommendation algorithm (i.e. 'if you liked film X, you may like films Y and Z')

that outperformed its own by 10%. The data were anonymised, and some of the data perturbed to inject noise (however, there was not very much noise in the data – too much noise would have made it hard for researchers to create and test a recommendation algorithm). However, two researchers based in the United States took a very short time to reidentify people in the dataset (Narayanan & Shmatikov 2008); the lack of noise in the system made this simpler, but their reidentification algorithm could have coped with far more noise than Netflix were able to add.

The following discussion of privacy appeared in the FAQs for the Netflix Prize.

Q. Is there any customer information in the dataset that should be kept private?

A. No, all customer identifying information has been removed; all that remains are ratings and dates. This follows our privacy policy ... Even if, for example, you knew all your own ratings and their dates you probably couldn't identify them reliably in the data because only a small sample was included (less than one tenth of our complete dataset) and that data was subject to perturbation. Of course, since you know all your own ratings that really isn't a privacy problem is it?

Narayanan and Shmatikov showed that identification from the dataset was remarkably straightforward, and that the amount of auxiliary information needed was not great.

Our conclusion is that very little auxiliary information is needed [to] de-anonymize an average subscriber record from the Netflix Prize dataset. With 8 movie ratings (of which 2 may be completely wrong) and dates that may have a 14-day error, 99% of records can be uniquely identified in the dataset. For 68%, *two* ratings and dates (with a 3-day error) are sufficient Even for the other 32%, the number of possible candidates is brought down dramatically. ...

Even without any dates, a substantial privacy breach occurs, especially when the auxiliary information consists of movies that are not blockbusters.

Where would the auxiliary information come from?

Given how little auxiliary information is needed to de-anonymize the average subscriber record from the Netflix Prize dataset, a determined adversary who targets a specific individual may not find it difficult to obtain such information, especially since it need not be precise. We emphasize that massive collection of data on thousands of subscribers is not the only or even the most important threat. A water-cooler conversation with an office colleague about her cinematographic likes and dislikes may yield enough information, especially if at least a few of the movies mentioned are outside the top 100 most rated Netflix movies. This information can also be gleaned from personal blogs, Google searches, and so on.

Narayanan and Shmatikov also exploited other online movie rating resources, such as the Internet Movie Database. The IMDb is a public arena for discussion of films, where people are happy to disclose their opinions (usually but not always under a pseudonym). It should be noted that IMDb discourages crawling its site, although an adversary would not be picky about sticking to its terms and conditions. However that may be:

Given a user's *public* IMDb ratings, which the user posted voluntarily to reveal *some* of his (or her; but we'll use the male pronoun without loss of generality) movie likes and dislikes, we discover *all* ratings that he entered *privately* into

the Netflix system. Why would someone who rates movies on IMDb—often under his or her real name—care about privacy of his Netflix ratings? Consider the information that we have been able to deduce by locating one of these users' entire movie viewing history in the Netflix Prize dataset and that *cannot* be deduced from his public IMDb ratings.

First, his political orientation may be revealed by his strong opinions about “Power and Terror: Noam Chomsky in Our Times” and “Fahrenheit 9/11,” and his religious views by his ratings on “Jesus of Nazareth” and “The Gospel of John.” Even though one should not make inferences solely from someone's movie preferences, in many workplaces and social settings opinions about movies with predominantly gay themes such as “Bent” and “Queer as folk” (both present and rated in this person's Netflix record) would be considered sensitive. In any case, it should be for the individual and not for Netflix to decide whether to reveal them publicly.

The analyses that Narayanan and Shmatikov anticipate are not guaranteed to be accurate. They are statistically-based, and so will have a built-in likelihood of error. They are also computationally very expensive. Nevertheless, they work.

Incredibly, however, Netflix persevered with its prize – won in 2009 by researchers at AT&T Labs – and because the prize was “such a research and business hit” (Lohr 2010) it immediately announced plans for another one. Only in March 2010 was the idea called off, in the face of privacy concerns, a lawsuit and the attentions of the Federal Trade Commission (Lohr 2010). Clearly the threats to privacy from deanonymisation remain poorly-understood outside the laboratory – perhaps partly because of the difficulty of the mathematics and statistics involved, but also because of a faith in the power of anonymisation that is beginning to look somewhat Pollyannaish. What the Netflix affair shows is that the simple distinction between identifying and non-identifying data is not sustainable, and should not be taken seriously by anyone with a genuine concern for privacy.

4.2.2 Techniques

The techniques used in this area trade on the fact that even supposedly non-identifying things about oneself are often unique or nearly unique. The clothes I wear today, the food I have eaten this week, the shopping I have done, the mileage I have driven, the places I have visited, the trains I have caught, the television I have watched – any of these mundane pieces of information could supply the key to my identity to someone who had access to sufficient background information. And if we are genuinely worried about privacy, then we cannot make the comforting assumption, especially in this over-wired world, that that background information will *not*, or only rarely, be available. Indeed, if a particular person was being targeted, the chances would rise dramatically that the information *would* be available, as blogs, tweets and Facebook pages are incredibly rich sources provided by data subjects themselves.

Not only that, but the quantity of information available simply grows and grows. The Web doesn't shrink (by much, anyway). We should assume that if information has once reached the Web, the chances are overwhelming that it will stay there, in some form or another, for quite some time. Hence even if the background information about some people was too small to be of value to an adversary at one point in time, we cannot assume that that will remain true over time. Conversely, if at any point there *is* sufficient background information about someone, that will remain so in the future.

The technique is very simple: given an anonymised database, and some auxiliary data, the adversary matches up lines which have attributes in common, as shown below. Suppose we have a line of data of the following form:

Attribute A	Attribute B	Attribute C	Attribute D	Attribute E
Value a	Value b	Value c	Value d	Value e

And a line of data of the following form comes into our possession:

Attribute C	Attribute D	Attribute E	Attribute F	Attribute G
Value c	Value d	Value e	Value f	Value g

If the conjunction of c, d and e is a quasi-identifier, then we can join the two lines of data to get the following amalgamated, and much more informative, line.

Attribute A	Attribute B	Attribute C	Attribute D	Attribute E	Attribute F	Attribute G
Value a	Value b	Value c	Value d	Value e	Value f	Value g

So, for example, in the Netflix data, if we take three obscure movies which some users have rated, and match them with the ratings given to those films by reviewers in the IMDb, then if we find a match between the two datasets, we can infer that the two lines of data are actually about the same person.

Using such techniques, databases can be linked together, joining them where a particular set of characteristics match. As noted, these techniques are hardly exact, and a join can only be made with a degree of probability that it is correct, but the inferences can be made with some measure of confidence (even if an inference is incorrect, this may not be of any comfort to someone who is falsely accused of doing something).

Data that are innocuous in isolation can be devastating in conjunction. For instance, suppose one database contained the following row:

Gender	Postcode	Children	Driver	Reference
F	SO17	2	No	Dr P. Mason

And another included:

Age	Sex	Address	Smoker	Registered GP
15	Female	Southampton	No	Dr P. Mason

Joining the two databases would allow us to create a more complete row of data which suddenly has more significance. If we could be reasonably sure that the same person was the subject of these two rows, a female resident of Southampton who has given the name of her GP as a referee in some transaction or another, that would tell us that the 15-year-old child of the second database was the mother-of-two of the first.

Another useful source of information is the structure of social networks. Narayanan and Shmatikov (2009) presented a framework for analysing privacy and anonymity in social networks with a re-identification algorithm targeting anonymised social network graphs (i.e. the basic structure of who follows whom, or who is friends with whom, without information about who these people are – information that is often shared). Their algorithm is robust to noise perturbation and other defences, and works even when the adversary doesn't have a great deal of auxiliary information. They

showed that a third of the users who can be verified to have accounts on both Twitter and Flickr can be re-identified in the anonymous Twitter graph with only a 12% error rate.

4.2.3 Impossibility theorems

The results presented intuitively in the previous section have been given mathematical expression in a number of powerful theorems showing the impossibility of achieving privacy while maintaining a given level of practical utility. The problem is that auxiliary information can always be brought to the calculation (cf. Dwork 2006).

Of course, to formulate mathematical theorems to show what is and what is not theoretically possible, mathematical definitions of the relevant concepts are required. One useful idea is that a database protects privacy if the adversary cannot learn anything about an individual from the database which he could not have learned without access (Dalenius 1977). Utility is very intuitively associated with the *unpredictability* of the output of the database as communicated to the user; if the user knew in advance what answers the database would give, then the database would not be very useful. The utility is also associated with the number of correct answers to queries that the database will give out; hence if we restrict the output of the database to protect privacy, its utility will be restricted proportionately.

If the adversary is allowed to bring auxiliary information to the table, then he can learn things that he should not be able to learn. If, say, the database is not supposed to give out personal data, but tells us something very bland and non-personal such as the average salary of a person working for a particular company, and the adversary knows that X works for that company and earns more than the average salary, then he is already able to narrow the range for X's salary. If he knows something more precise about X (say that she earns 50% more than the average salary), then he can calculate her salary more accurately. Indeed, suppose X did not work for the company, but that the adversary knew that she earned 50% more than the average salary for that company; in that case the adversary could work out X's salary *even though she does not appear in the database at all*.

This is of course an intuitive expression of what turns out to be a fairly complex mathematical idea, but a formal, mathematical expression can be found in (Dwork 2006) and elsewhere.

4.2.4 The legal repercussions

The legal repercussions of this work could be very large. Ohm has argued that it makes the EU directive far too wide for any practical purpose. Virtually anything will be illegal because

... anonymization makes laws like the EU Data Protection Directive overbroad, in fact essentially boundless. Because the Directive turns on whether information is "directly or indirectly" linked to a person, each successful reidentification of a supposedly anonymized database extends the regulation to cover that database. As reidentification science advances, it expands the EU Directive like an ideal gas to fit the shape of its container. A law that was meant to have limits is rendered limitless. A careful balance struck by legislators between privacy and information flow shifts wildly to impose data handling requirements to all data in all situations. (Ohm 2010)

In other words, the directive applies to everything, because someone could be identified from any piece of data of which he is the subject, even if anonymised. Nevertheless, powerful though Ohm's reasoning is, I am not convinced that this sweeping conclusion is necessarily the case – there is I would suggest some leeway in the directive's phrasing: "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". As Ohm points out, indirect identification is included, but there is a question about whether, say, one's film preferences, as recorded by Netflix, or even one's search query terms as recorded by AOL, count as matters specific to someone's identity in the words of the directive, at the point of data release.

Furthermore, the Data Protection Act 1998, which implements the EU Directive, is much narrower, referring to personal data as "data which relate to a living individual who can be identified from those data and other information which is in the possession of, or is likely to come into the possession of, the data controller." Hence if we take Ohm's arguments about deanonymisation and the directive on board (and cf. also Crossman et al 2007, 120-121), it begins to look like the EU Directive forbids the release of almost everything, while the UK Act, which supposedly implements it, will allow rampant reidentification!

Thirdly, the deanonymisation issue means that virtually everything is tagged as potentially dangerous. Suppose a number of datasets need to be plundered before a particular individual X is identified with a reasonable level of certainty. How does the law stand with respect to these? Are all the data releases culpable? Or is it just the last one, which was, as it were, the straw which broke the identity camel's back? After all, all the other datasets were released into a world in which the auxiliary information required to make the identification was not generally available. Or was it the last dataset which came into the possession of the adversary, which may not have been the last one released? These really do appear to be open questions.

I think it more likely that UK law will allow data releases that Ohm would believe to be reckless. The issue then is where to draw the line – and in my view, given the importance of technical advances here, it will be very hard to encode a sensible view of what is a 'safe' release of data into either technical, computational, statistical or legal instruments.

4.2.5 Responsibility

If Ohm's argument, that total responsibility must be distributed widely, doesn't stand, it is important to thrash out where responsibility for a data breach would lie in the event that a jigsaw identification did occur. There are a number of possibilities, none of which is very satisfactory.

Perhaps the most obvious candidate would be the provider of the final piece of the jigsaw that enabled the identification to take place. Yet this surely suggests an extremely onerous duty on anyone releasing anonymised data into the public domain, without powers to control access: that they should ensure that no-one can be identified, in tandem with the information context, from the data. This would be a Sisyphean task, and would have a terribly chilling effect on transparency. Furthermore, it would be very unfortunate if, say, indiscreet postings on social networking sites by data subjects were to prevent the orderly release of valuable government datasets, because the probability of identification of those subjects, given

the auxiliary information they themselves had made available on the social network, was too high.

Yet the order in which the pieces of the jigsaw are made available does seem important, even when some of the most telling pieces are released early on in the process. This is a particularly complex issue to resolve, especially as there are two relevant classes of responsibility, the legal and the political, to be taken account of. It may be that the sheer complexity of the analysis required to deanonymise will protect those hoping to release valuable datasets, by rendering deanonymisation of any carefully-anonymised particular dataset unlikely.

4.2.6 Empirical evidence

Empirical evidence is lacking in this area. Many issues of data protection turn upon the effort required for a privacy breach, and the likelihood of that occurring. We know that deanonymisation is a real possibility – mathematically from the impossibility theorems, and in practice from the deanonymisation of the AOL data, performed, in Ohm's phrase, by "a small group of bored bloggers." Ohm makes the point that this is not just a theoretical possibility; we cannot simply ignore the issue on the ground that only a super-adversary could possibly marshal enough computing skill.

It is certainly true that deanonymisation requires a lot of skill and resources. However, the costs can be amortized or written off over time as it needs only be done once and yet could provide sensitive details of thousands or even millions of individuals (Narayanan & Shmatikov 2010), while the more information that is revealed about a person, either via legitimate data releases or by deanonymisation of anonymised databases, the easier it is to expose that person elsewhere in the future (Ohm 2010).

Furthermore, Moore's Law means that more computing power is available over time (the amount of power per chip doubles every 18 months). If we map this technology development cycle onto Britain's electoral cycle, this means that a technology that is feasible at the beginning of a Parliament, and whose benefits are monetisable, will be in routine use by the end of it.

Nevertheless, we do lack good empirical evidence about how data will be used in the real world. Bringing databases built on different principles together is a hard problem (Garfinkel 2008, Cavoukian & El Emam 2011), and data scraped off the Web by criminally-minded adversaries might well be harder to deal with than the well-structured datasets released by AOL and Netflix. Furthermore, there may also be a difference in capabilities between well-funded and highly-qualified academic researchers and their teams of assistants, and the resources available to real-world adversaries. The small number of empirical studies there have been seem to confirm the difficulty of the problem (Lafky 2009), but the work must be ongoing because of the advantage that the continuing relevance of Moore's Law gives to the adversary. It is important to address this lack of evidence somehow, a point to which I shall return in my recommendations below.

4.2.7 Data integrity

Recall from section 3.1.4 the importance of retaining the integrity of data through any transparency procedure; if releasing data caused problems in the data, or made the use of the data more problematic, then a rethink may be needed. In certain circumstances, the release of anonymised data could present a problem for legitimate data users.

For instance, it may be that a publicly available dataset could be used to identify data subjects in combination with other datasets that are available to accredited researchers under controlled conditions. Of course, this would not present a serious *privacy* problem, because government agencies have a range of constraints and sanctions that they can apply, if need be, to trusted and accredited researchers (and no-one else would have access). Hence data subjects are extremely unlikely to be identified via this route, because there will still be several methods in place to prevent or deter identification.

However, if data subjects remain *identifiable*, the problem *for researchers* may remain, because their access to the restricted datasets could be affected by a new-found and unwanted ability to identify some data subjects. It may also present a problem *for providers*, as (particularly with sensitive data such as health or education data) they may find themselves with more stringent requirements under the Data Protection Act. Even suppliers of raw data (such as GPs or schools) may be more reluctant to continue to provide such data where they are under no statutory obligation to do so, if they begin to judge that the burden of their responsibilities and liabilities is higher than they wish to bear.

The solution to this issue is easy enough; engage with data users to ensure that anonymisation of publicly-available datasets is sufficient to maintain the anonymity of data subjects across the range of data releases, including those datasets to which access is restricted and controlled. Both data providers *and* data users need to be satisfied that no-one is identifiable from the data provided.

4.2.8 Let's not say 'anonymise'

It is essential that policymakers, data managers, data controllers, privacy officers and lawyers do not automatically assume that 'anonymised' data cannot be used to reidentify people. This is a semantic and legal fiction which could not have been propagated without the estrangement between legal and technical practitioners which I discussed in section 1.2.4. As noted there, *legal discourse alone is not sufficient to address this problem*.

Following Sweeney, Ohm suggests, not unreasonably, that the term 'anonymise' be removed from the legal vocabulary. Its use implies that data have been placed in a state wherein an individual cannot be identified even after some manipulation. This is now in some dispute to say the least, and so the term – an achievement word – should be replaced by one that connotes effort without necessarily guaranteed success. Sweeney likes 'deidentify', Ohm prefers 'scrub'. My own favourite, for what it is worth, is 'disguise'. Someone can attempt to conceal their identity with a disguise, but that does not imply that they cannot be identified despite the disguise; Superman disguises himself as Clark Kent, but his identity is still blindingly obvious to everyone except the characters in the show. At the moment there is little consensus about this, so in this report I have and will continue to use the unsatisfactory terms 'anonymise' and 'deanonymise'.

The key point to emphasise is that there are no cast-iron guarantees in this space, and we shouldn't continue to pretend there are. One plausible reading of the Directive ("one who can be identified") is that without such guarantees we have to consider data to be personal data. As (Cavoukian & El Emam 2011) point out in the context of health data, that would have a chilling effect on the sharing of valuable data.

And it should be emphasised that just because anonymisation techniques are susceptible to jigsaw identification, it does not follow that we should *stop* anonymising (Cavoukian & El Emam 2011), any more than we should stop using burglar alarms to protect our homes even though they are ineffective against the most skilled burglars. Even if they are susceptible to attack, they make the task of identifying individuals very much harder. And if the sharing of anonymised data became as legally and administratively burdensome as the sharing of personal data now is, then the incentives for anonymising would evaporate. However, the possibility that data can be deanonymised, combined with the great uncertainty about how difficult a problem that would be, does need to be factored into our discussion about whether data can be released into the public arena.

4.3 Alternative approaches

It is important to recognise that it would be a terrible defeat if it became impossible to publish useful datasets because of deanonymisation. Commentators on this issue do have potential solutions to the problem, but there is a problem in squaring them with the demands of a transparency programme. I shall discuss these matters in this section.

4.3.1 Disclosure control

The first broad class of methods to protect privacy is what is called 'disclosure control' (cf. Anderson 2008, 275-311). Under disclosure control, also known as query control, the data controller interacts with those querying the dataset, and prevents privacy-threatening queries being answered. As Narayanan and Shmatikov point out, "this can be a hard pill to swallow, because [disclosure control] requires designing a programming interface for queries, budgeting for server resources, performing regular audits, and so forth" (2010, 26).

Broadly speaking, the idea of disclosure control is to screen queries of the data for a set of queries that between them will reveal a sensitive statistic. For instance, if I want to know X's salary, I might query a database for the average salary of everyone in the database, and query for the average salary of everyone but X. Guarding against this can become computationally very intensive.

The problem with disclosure control in the context of privacy is obvious: the philosophy of transparency views data as public property, and advocates open, free and unrestricted use of released data, including allowing passing those data to other people. Furthermore, one of the aims of the UK transparency programme is to avoid a heavy bureaucratic overhead in order to facilitate the publication of as much data as possible without departments or agencies incurring large costs. Hence the whole idea of disclosure control as a routine means of protecting privacy is simply incompatible with transparency. The lower the overhead that disclosure control imposes on the data release process, the better, but in a context where the aim is to allow citizens to download data on open licences with as few obstacles as possible, then disclosure control is clearly very contrary to the spirit of transparency. Indeed a tightly controlling mechanism is relatively difficult to implement whenever a large number of queries is anticipated; a tough disclosure control regime would be something of a bottleneck.

That does not mean of course that transparency is necessarily threatening to privacy. It does mean that the use of one powerful but expensive mechanism is denied to any transparency programme in which privacy is respected.

Compare the situation with transparency as defined here, with that described by (Cavoukian & El Emam 2011), who argue for the continued practice of anonymising and sharing health data. Cavoukian and El Emam's assumption that deanonymisation will be hard is boosted by the fact that health data will only be shared with vetted and trusted third parties. This means that the third parties are accountable, and the data sharer can be very much more confident that adversaries prepared and equipped to deanonymise will not get their hands on the data. This strengthens Cavoukian's case, but her powerful argument depends on supplementing anonymisation with disclosure controls. This is not an option for a transparency programme in the scope of this report.

4.3.2 Administrative mechanisms

Similar arguments apply to the use of administrative mechanisms such as consent, terms and conditions, registration and charging. These do have lower bureaucratic overheads, and would allow the government to get around data protection constraints – if consent is gained for the release of data by data subjects, then data controllers have a much freer hand, even if (as noted in section 3.1.5) consent at the moment is a very unsatisfactory mechanism from the point of view of empowering the citizen.

Consent means getting the agreement of the data subject for the data to be used for a wider set of purposes than the immediate purpose of collection. With terms and conditions, the person downloading the government data from the Web (the user) would have to tick a box to agree to a set of terms, which could include, e.g. restrictions on the uses to which the data are put, constraints on who the user can pass the data onto, or a prohibition on the user allowing the data to be transferred, either by himself or by a third party, outside the OECD (or jurisdictions where data protection law was well-policed and governed by international agreements).

A registration system would require a user to register and identify himself before getting access to the data. Then if data were misused, there would be at least a possibility that he could be traced and sanctions placed upon him (the deterrent effect may only be small, but would at least be non-zero). In the event that a particular set of data was released on a regular basis (such as the crime data, for example, which at the time of writing are released monthly), then the government could run a regular feed service where registered users could receive the up-to-date data automatically upon their release. If the government then discovered that information was being misused by a particular user, then he could be deleted from the regular feed, or from the registration entirely.

A charging system is obvious: the government simply charges a small fee for access to the data, which would add to the costs of those planning to use the data to compromise privacy, and would alter the economics of any criminal or anti-social usage.

These mechanisms are crude and not too easy to enforce, and do impose a bureaucratic cost, both on the government to run schemes like this, and on the users who would have to comply with them. They also raise the barriers to entry for new information entrepreneurs to enter the market. The quality of the data could also be compromised in the event that, say, only a few data subjects gave informed consent.

Hence they are not very compatible with the transparency programme, and should not be viewed as a useful general purpose tool in this space.

However, it should be remarked that the costs of such schemes are not enormous. It might be the case that in some circumstances, the government could consider using such schemes for potentially sensitive data that it wished to make available. I did suggest in my interim report on the crime data that exploration of a consent-based system, in which some victims of crime could agree to having geographical data about the crimes of which they were victim being given without being made deliberately vague, would be a good idea – although, because of the bureaucratic overhead involved, I did not recommend that such a system should definitely be used.

One could imagine their use in some circumstances, including:

- Where there were some identifiable risks from releasing the data.
- Where the data were not obviously sensitive in the abstract, so the opportunity from a data release was of greater moment than the threat. Sensitivity of course depends a great deal on context, and the use of an informed consent mechanism would allow the data subject a say on data that might not otherwise be released.
- Where government departments and agencies could be confident that the costs to them would be made up by wider economic benefits.
- Where there was keenly-expressed demand and a willing user base for the data.
- Where there was evidence that there would be many information entrepreneurs willing to compete against each other to provide information services using the data, so it was known that the barriers to entry to the market were not high enough to deplete competition.
- Where the generation of the data by the government department or agency demanded substantial investment, thereby increasing the desirability of creating an income stream.

In short, bureaucratic access control systems can be of value when (a) having them in place makes a data release possible where it would not otherwise be possible (which might be for privacy reasons, or other reasons such as cost), or when (b) their presence would have a positive effect on the quality of the data being released.

In such circumstances, then one could imagine one or more of these methods having some value. However, at the moment, there would seem to be no particular reason to explore these ideas, as there are plenty of datasets of very little risk indeed which could be released without compromising privacy at all. They are worth bearing in mind, however, for future releases, especially as the transparency programme becomes more demand-driven.

Hence access control systems are unlikely to be part of the model of transparency in the general case, but they could add value in certain areas, in particular by allowing some privacy concerns to be allayed. Which of the methods – consent, terms and conditions, registration and charging – would be suitable in what combination would depend on the sustainability model for a particular data release. Who should bear the costs? Is the demand for the data sufficiently great that cost is not an issue? Who should administer the bureaucratic workload? Is access control compatible with low

barriers to entry? How should competition be fostered? And of course what is the threat, if any, to privacy? All of these issues would need to be addressed in each particular case.

4.3.3 Differential privacy

Recall the impossibility theorem of section 4.2.3. It says, in effect, that the possibility that something can be learned about one from a database cannot be reduced to zero. This conundrum has led to an alternative mathematical definition of privacy, which states, broadly speaking, that the risk to one's privacy should not substantially increase as a result of one's appearing in a statistical database. This notion has been called *differential privacy* (Dwork 2008).

Note that differential privacy is a relative, non-absolute notion. There are no cast-iron guarantees – this informal characterisation relies on notions of risk, which is not to be 'substantially' increased. Of course, what counts as a 'substantial' risk may differ depending on whose beholder the eye belongs to, on changing social mores, and on the public good of increasing the utility of a released database. Nevertheless, the privacy guarantee from differential privacy is still very strong because it is a statistical property; once the appropriate level of risk has been decided on (a social question), the risk can be maintained independently of both the computational power and the auxiliary information that the adversary is able to throw at the problem. In effect, it is a concept that judges the *computation* as privacy-preserving or otherwise, rather than trying to make an impossible distinction between identifying and non-identifying *data*. It also implies that, even if a privacy breach does occur, the data subject can be assured that it was not the presence of her data on the database that caused the problem.

Differential privacy makes an important contribution to the field. However, once more a problem with it is that, so far as we know, to be exploited, it requires the controller of a dataset to retain control and to administer the access granted to outsiders. Control using this methodology involves monitoring queries via a special language, the Privacy Integrated Query Language (PINK – McSherry 2010), to produce a precise statement of the privacy that has been revealed as a result of its associated theory quantifying privacy revelation. The database itself must be protected in order to monitor queries.

An important advantage of the exploitation of differential privacy is that it allows setting a level of privacy compromise which is tolerable. This of course is a political decision, not a technocratic decision, to make it larger or smaller. However, once it has been set, the use of PINK allows a precise specification of how far a set of queries has approached the limit.

The major disadvantage for the method is that, once the limit has been reached, the dataset is finished; once the privacy 'budget' has been expended, the dataset cannot be exploited further without threatening privacy beyond the limits originally set. Of course, in such circumstances, there will also be a political threat that those wanting to extract further utility from the data will apply pressure to raise the privacy limit as it is approached. To avert this the controller of a dataset must retain control and must be resistant to pressure from those wishing to exploit the data beyond the limit (who may, of course, include the controller's bosses).

However, that is not to say that differential privacy could not have a role in the context of transparency. Differential privacy is underpinned by a theory that (a) is

rigorous, and (b) defines 'non-identifying' on the basis of what might be done with the data, in addition to the properties of the data. This means that it might well have an important role to play in estimating the dangers to privacy of a particular data release, and giving what we might call a 'privacy baseline' against which to judge the feasibility of release.

The theory may also become more practicable as it becomes linked to other notions of privacy. For example, an early result has appeared in ArXiv (Li et al 2011) that in effect analyses the flaws in k-anonymity, and then supplements it with a more relaxed notion of differential privacy to provide improved privacy guarantees from 'hiding in a crowd of k people'. It is too early to tell whether this approach is the way forward in this area, but the paper certainly indicates that privacy and transparency advocates should watch this space over the next few years.

4.4 How should we deal with anonymised data?

Opinions about how we should move forward differ very markedly. In this final section, I shall comment broadly on the arguments about anonymisation and other data safety mechanisms, and will try to suggest ways to assess and properly manage the risks involved.

4.4.1 The low risk of deanonymisation

There are clearly risks with the use of anonymised data, but they are relatively low and the various potential solutions to the problem are often too risk averse (for instance, using differential privacy will prevent data being used by all but very trusted and accredited researchers). We should beware of throwing the transparent baby out with the private bathwater. After all, anonymised data have been published for years, even in sensitive areas. Raiseonline (<https://www.raiseonline.org/Login.aspx>) provides interactive analysis of school and pupil performance data, extracted from the National Pupil Database, while the Secondary Uses Service (SUS – <http://www.connectingforhealth.nhs.uk/systemsandservices/sus>) provides comprehensive data on health service outcomes. Services like these have served vital public purposes, and also have helped generate best practice in this area, although SUS in particular has occasionally been the subject of some controversy (e.g. Brown et al 2010).

In an interesting riposte to Ohm, Ann Cavoukian, the Information and Privacy Commissioner of Ontario, has argued that "the claim that the de-identification of personal data has no value and does not protect privacy due to the ease of re-identification is a myth" (Cavoukian & El Emam 2011).

Cavoukian and El Emam base their argument around three claims that are important correctives to the argument made by Ohm. First, they point out that anonymisation remains a strong tool. It does protect privacy, of course – anonymised data are certainly more private than the personal data they replace. The steps to deanonymise are not at all trivial; so far most of the more sensational results have been demonstrated by highly motivated, well-resourced and well-trained researchers. Hence the anonymisation of personal data clearly puts a barrier between the data subject and those who would misuse the data about them. As an analogy, just because some burglars are capable of neutralising burglar alarms, that does not mean that we should stop using them to protect our houses. Their protective value may not be 100%, but it still outweighs their cost. Similarly with anonymisation.

Secondly, they argue that deanonymisation is not just a tricky step for the adversary; it is an enormous barrier. The empirical evidence, such as it is, shows that even experts struggle to deanonymise datasets. It is also a very costly procedure, which would be beyond the reach of many.

Thirdly, they point out that if the argument of (Ohm 2010) is accepted without reservation, then in effect we would have to treat anonymised data as identifiable, and therefore as personal data. Hence the administrative overheads for someone dealing with anonymised data would be as extensive as the overheads applying to people dealing with personal data (e.g. the need to contact all data subjects every time the data were to be used for a reason other than that for which they were collected). This would be an intolerable situation, as there would then be no incentive for anyone to anonymise anything at all. That would be crippling either to privacy, or to the benefits of data sharing. The consequences of Ohm's argument go way beyond the potential harms he has identified.

A fourth point which they could have added is that anonymisation techniques are becoming increasingly sophisticated. In contrast, the AOL data were anonymised using pretty basic techniques. The 'doomsday scenarios' of widespread privacy breaches as a result of anonymised data sharing have not happened to those datasets which were anonymised more carefully (indeed, even the most notorious losses of unencrypted personal data have not led to any serious outcome – Kelsey 2009). The risk may be non-zero, but it may also be close enough to zero to render it effectively negligible. In that case, it would hardly be a sufficient ground to deny ourselves the benefits of sharing data and transparency.

4.4.2 Limits to the optimistic argument

This is an important corrective to the view of Ohm and others. Cavoukian's cogent case that anonymisation still has a vital role to play in the preservation of both privacy and data-value is absolutely right. However, in the context of transparency as defined in this report, it is also important to understand the limits of her argument. She focuses on one particular area of highly sensitive information, personal health information, which is very private and has great social value. In this context she is thinking mainly of secondary uses such as "research and evaluation", and other "authorized secondary purposes". The anonymised data are most likely to fall into unauthorised hands by accident or via an opportunistic "inside job" and so "it is unlikely that the person who finds the information would have the motive or capacity to attempt to re-identify the individuals in the data set".

Absolutely. However, Cavoukian's argument is not conclusive. The context upon which she focuses allows a much greater level of control than the transparency with which we are concerned in this report. The assumptions she makes do not hold in the uncontrolled online context. The data, once online, will remain online, and cannot be withdrawn (the data released by AOL in 2006, for example, were withdrawn within three days when AOL realised its mistake, and yet they are still widely available). It is, as she says, no big deal if the incompetent or the opportunist can get hold of anonymised data, but once they are online the data can be sought out by competent and well-resourced adversaries at their leisure. It does not matter if such masterminds are rare; they will come to the data, rather than waiting for opportunity to knock.

Cavoukian makes much of the "conclusion that the re-identification of individuals is a difficult and time-consuming task, on the part of skilled technicians." She is once

more absolutely correct to do so. However, this is a very grey area – we simply do not know exactly how hard a task this is. There is little empirical evidence. But in the context that interests us, time is not necessarily a problem for the adversary, because the data cannot be withdrawn from circulation if they get into unsavoury hands. Neither is difficulty necessarily a problem, because skilled researchers can be recruited and brought to the data on a schedule that suits the adversary. Cavoukian's argument assumes a window of opportunity that can be closed if there is an unacceptable risk. This is not the case if datasets are published online.

This consideration also undermines another foundation of Cavoukian's argument, that the costs of deanonymisation are too great for most adversaries. Again this is probably true, though empirical evidence once more is lacking (and organised criminals do have access to impressive resources). But in the context of transparency as defined here, the costs of deanonymisation can be amortised – that is, written off against the potential benefits stemming from the large number of identities that may be uncovered – and distributed, by selling information on in the many criminal markets for stolen identities.

Furthermore, the science of deanonymisation will no doubt develop, as will the cheap computing power available to an adversary. Even if an anonymised dataset is uncrackable now, it may not be two years hence – and, even if it has been officially withdrawn by then, it must be assumed that its Web presence will continue.

Finally, deanonymisation gets easier as the amount of auxiliary information increases. And of course relevant auxiliary information is increasing online all the time. In particular, if a government is releasing anonymised data in a sector regularly, then the cumulative effect will be to increase the amount of relevant data in that sector over time. Each release is not only vulnerable in itself, but it will act as potential auxiliary data to crack other datasets in that area.

4.4.3 The need to confirm our optimistic intuitions

Cavoukian's closing statement that “while de-identification may not be a perfect solution to reduce all privacy risks when personal information is being considered for secondary purposes, it is an important first step that should be used as part of an overall risk assessment framework” (Cavoukian & El Emam 2011) correctly draws attention both to the value and the limitations of anonymisation. The task of any transparency programme such as that of the United Kingdom is to craft the risk assessment procedure in a realistic and conscientious way. I shall address that task in the recommendations below.

The risk of deanonymisation is very low, but not zero. However, even critics of anonymisation admit that the situation is extremely unclear; for example, Narayanan has written that “we need to better understand the theoretical limits of anonymization and to extract the common principles underlying the more complex re-identification techniques developed in recent years” (Narayanan 2009). There are a few examples of dramatic deanonymisations in the literature, some of which are reviewed above, but they have tended to be of well-structured datasets under optimal conditions. This is a tiny percentage of all the anonymised data releases over the years, and the heightened awareness of the potential for deanonymisation has not led to a flood of deanonymised data.

It is also important that debates about particular anonymised datasets are properly informed, and that (recalling section 3.4) we are transparent about transparency

practice. In order to maximise the value of debate about releasing anonymised datasets, it will be important to be open about which anonymisation techniques are being used to anonymise datasets (this will not allow the original dataset to be reverse engineered, of course), in order to allow realistic and independent assessments about risk to be made.

All this is cause for optimism that sophisticated anonymisation, perturbation and pseudonymisation techniques will continue to allow the release of valuable data for use by the public, and the management of a negligible risk. However, it is important to confirm that intuition with further investigative work to show that the risk really is negligible – and if it proves not to be, to suggest further ways forward. It is not the place of this report to 'solve' this problem or 'decide' the argument one way or the other; my aim here is to state the position of the debate as it stands, and suggest means by which realistic consideration of the issues surrounding particular sensitive dataset releases can be facilitated.

5 Conclusions and recommendations

5.1 Introduction

I have reviewed transparency, privacy in the context of transparency, and the technological developments that have created so many imponderable problems. The general outline of my recommended courses of action is probably clear from the discussion; we need better institutions and conversations to screen data for the privacy implications of their release, and we need to include technologists in these conversations to a much greater degree than has historically been the case.

To this end, I make fourteen recommendations. None of them is intended to place a large bureaucratic overhead on the UK government's transparency programme, nor should there be any substantial cost implications (certainly not relative to the potential gains across the economy by improved transparency). All the recommendations are supported by the reflections above. Some will require ongoing effort, others could be implemented immediately.

There have been some ideas floated in the text above which I have not included in the recommendations (for instance, my hope that the field of differential privacy might in time contribute to our understanding of privacy, in section 4.3.3, or that a practical architecture for citizens' consent management may emerge from British research, in section 3.1.5). That is not to say that these ideas should be discounted or ignored, but they rest on the potential of currently progressing research rather than on proven principles or easily implemented systems, and are considerations for the longer term only.

5.2 Conclusions

The outlines of the conclusions should be clear by now.

- Privacy is extremely important to transparency. The political legitimacy of a transparency programme will depend crucially on its ability to retain public confidence. Privacy protection should therefore be embedded in any transparency programme, rather than bolted on as an afterthought.
- Privacy and transparency are compatible, as long as the former is carefully protected and considered at every stage.
- Under the current transparency regime, in which public data is specifically understood not to include personal data, most data releases will not raise privacy concerns. However, some will, especially as we move toward a more demand-driven scheme.
- Discussion about deanonymisation has been driven largely by legal considerations, with a consequent neglect of the input of the technical community.
- There are no complete legal or technical fixes to the deanonymisation problem. We should continue to anonymise sensitive data, being initially cautious about releasing such data under the Open Government Licence while we continue to take steps to manage and research the risks of deanonymisation. Further investigation to determine the level of risk would be very welcome.

- There should be a focus on procedures to output an auditable debate trail. Transparency about transparency – metatransparency – is essential for preserving trust and confidence.

In the remainder of the report, recommendations will be made which are intended to implement these conclusions without making too strong a claim on scarce resources.

5.3 Recommendations

There are fourteen recommendations. Some are quite general, while others suggest specific actions to be carried out. Although any or all of them could be adopted in isolation, they are mutually supportive, and are intended to work together as a package. They are not intended to place an excessive administrative or budgetary burden on government. They should allow the transparency programme to progress while preserving the confidence of the British public. The ideas are intended to appeal across party political boundaries, and to parties in both the Coalition government and the opposition. I also hope that these recommendations will keep the UK transparency programme to remain in the vanguard of innovation in this area, while also helping it learn from the positive experiences of other governments.

Recommendation 1: Represent privacy interests on the Transparency Board

It is vital to preserve public trust and confidence in the transparency programme. To that end, as it will prove impossible for the programme to avoid questions about personal data, it will be far better to embed privacy protection in the programme itself, rather than as a bolt-on or additional component of the procedures.

The obvious way to do this from the top is to include someone independent of government in the Public Sector Transparency Board advisory body whose role will be specifically to protect citizens' privacy.

Such a person should be recruited to the Transparency Board with an international reputation in privacy advocacy, particularly someone with a clear understanding of the complex technical and technological issues. The new recruit should be able to command the confidence of those concerned with the protection of privacy, and be of sufficient stature to defend privacy interests effectively in the Transparency Board.

One candidate for this would be the Information Commissioner, whose public role is precisely that. Furthermore, the IC also has responsibility for promoting freedom of information; hence he or she should not be inclined simply to act as a blocking force. However, it may be that the IC's role as independent of government would preclude him or her from this role. It may be that the IC should be seen as independent of, and therefore free to criticise, the Transparency Board. That is a reasonable view of the IC's function, in which case the new recruit for the Board should be sought from the wider community.

Recommendation 2: Use disclosure, query and access controls selectively

There is a potential clash between the transparency agenda, and the increasing technical consensus (Narayanan & Shmatikov 2010, Ohm 2010) that disclosure, query and/or access control measures will be required to allow anonymised data to be released online.

Such measures would be detrimental to transparency. I have therefore not attempted to suggest that they should be used routinely in the transparency agenda. My recommendation is not that disclosure controls are useless or too expensive, but merely that their routine use cannot be a part of the transparency programme as currently conceived.

This should *not* be taken to mean that personal data can be released without controls. It *should* be taken to mean that: *if* a data release were potentially privacy-threatening, *and if* disclosure/access controls could remove the threat, *then* they should be considered. *If* resources did not permit the implementation of such controls, *then* the data could not be released.

Hence controls are not ruled out, and could be used in certain circumstances (cf. section 4.3.2). The particular measures that are likely to have a role are:

- Consent.
- Use of terms and conditions.
- Use of registration to identify users.
- Charging.

The circumstances where their use could be valuable include:

- Where there are identifiable risks from releasing the data.
- Where the data are not overtly sensitive in themselves.
- Where government departments and agencies could be confident that the costs to them would be made up by wider economic benefits.
- Where there is keenly-expressed demand and a willing user base for the data.
- Where the barriers to entry to the market are not high enough to deplete competition.
- Where the generation of the data by the government department or agency demands substantial investment.

Control systems should be considered when:

- Having them in place makes a data release possible where it would not otherwise be possible.
- Their presence would have a positive effect on the quality of the data being released.

Not all the relevant considerations here are concerned with privacy, though they will cover a number of privacy-threatening situations.

Standard methods for protecting privacy such as anonymisation should continue to be used, even though they cannot give a 100% guarantee that they cannot be undone by a sufficiently adept adversary.

Recommendation 3: Include the technical paradigm

Legal definitions of privacy have tended to dominate the debate in the United Kingdom and elsewhere. However, these have proved inadequate to provide a clear framework for analysis of privacy issues, especially in the context of jigsaw identification using recently developed deanonymisation techniques.

To this end, there should be greater awareness of the technological paradigm. This should happen in two specific ways.

1. Technologically-trained experts should be brought into procedures for deciding whether or not to release particular datasets. A description of how this may be done is given in recommendation 7.
2. There needs to be a greater awareness of technical issues in the Information Commissioner's Office (ICO). The ICO has made welcome strides in recent months, for example with the appointment of a Principal Policy Advisor in this area, and the creation of a Technology Reference Panel. Nevertheless, the severe technical demands made by cutting-edge research in deanonymisation mean that more effort is needed in this direction.

Recommendation 4: Move toward a demand-driven regime

The transparency programme covers two separate agendas as noted in section 2.1.4, an accountability agenda and an information agenda. In its current phase (a year or so into the Coalition government), it has concentrated on accountability, and as such has something of a 'top down' feel. Ultimately, it would be desirable to move the emphasis from accountability to information, providing the raw materials for citizens, charities, intermediaries and entrepreneurs to develop a rich picture of their communities, to enable and empower users to interact more effectively with their fellow citizens, organisations, companies and government. As this happens, the transparency programme should shift to a more demand-driven, 'bottom-up' regime.

The Transparency Board is attempting to serve both the accountability agenda and the information agenda, but (notwithstanding the presence on the Board of a successful information entrepreneur and other advisors who have experience in the use of data by citizens) attempts to suggest which datasets might be useful for citizens in real-world contexts are naturally somewhat hypothetical.

In a demand-driven regime, information entrepreneurs would ask for the datasets they felt they needed, or felt that they could use to create value, whether social value or commercial value (profit) for their own firms. This suggests two requirements.

1. Entrepreneurs must know what datasets there are.
2. There must be a screening process to ensure that privacy-threatening releases (and other problematic releases, such as ones which might threaten national security) could be challenged and blocked.

These two requirements will be the subject of the next two recommendations.

A demand-driven regime would, as argued in section 2.2.2, pose some threats to privacy that are not currently on the horizon in the present context. These threats are not terminal, and should be addressed using the procedure outlined in recommendation 7.

However, it is worth noting here that in two respects a demand-driven regime would *promote* privacy. First, it would be incumbent on those demanding the data to demonstrate conclusively that it was either not privacy-threatening to release them, or that their release, and the use to which they would be put, were of overwhelming public interest and proportionate compared to the privacy threat. If such a case was robustly expressed and rigorously scrutinised, this would be a good indicator of the likely threat to privacy of that particular release. This compares favourably to the current situation, where the release of data is uncontextualised with a small understanding of the demand side. As noted in section 3.1.4, privacy norms and expectations are highly context-dependent, and very difficult to state convincingly in the abstract.

Second, two important principles, the Purpose Specification Principle, and the Use Limitation Principle (cf. section 3.1.2) are in tension with the transparency programme's driving assumptions that *serendipitous reuse* of data (i.e. gaining value from data by its reuse by others in unanticipated contexts and for unanticipated purposes) should be facilitated, and that productive use of data cannot be fully anticipated by data controllers. The idea of serendipitous reuse does not allow those releasing data to specify exactly how they will be used, while the demands of the transparency programme for a small bureaucratic overhead make terms and conditions

and other access control methods problematic (section 4.3.2 and recommendation 2). However, if a case has to be made by someone *demanding* the data for a specific purpose, then at least those judging threats to privacy will be in a position to understand *some* likely contexts for their use.

Recommendation 5: Create a data asset register

A register of government data assets should be compiled and publicised. This need not be complete (indeed, could not be). The register should set out what datasets were controlled, what they contained, and what decisions had been taken about their release. Possible classifications would include:

- The data are confidential and on no account will be released via the transparency programme.
- The data are accessible without restriction from the Web on an OGL.
- The data are accessible from the Web with some restrictions (if access control has been deemed useful in this particular case, or if for some reason an older, more restrictive licence is in force).
- The data are not deemed confidential, and though not currently accessible from the Web, they are scheduled for release in a named format at an appointed time.
- The data are not deemed confidential, and are not currently accessible from the Web or currently scheduled for release, in which case a request for them can be lodged and processed.

The asset register could be centrally curated, or kept by individual departments and agencies.

In the case of confidential data, it is important for metatransparency (section 3.4) that citizens are aware of what information government holds even if it is not made available. When citizens use public data to build pictures of their communities, they need to know whether such pictures are complete, or whether certain aspects are under-represented. There is nothing necessarily sinister in keeping data confidential, but the fact of the restriction should be made clear (so it could be challenged in public debate, by an official such as the IC, or even in the courts).

Note also that the creation of a register, and its use by information entrepreneurs, is likely to improve the quality of government data, by providing feedback to influence collection methods, ontological assumptions, quality, reliability, timeliness, output formats and so on.

Given the register, entrepreneurs should be able to ask for particular datasets which were currently unavailable. Note:

- This would not entitle the requester to exclusive use of the data. The presumption would be that the data requested would be placed on the Web via some access point such as data.gov.uk.
- If data of a particular type were requested, the request should be expanded to cover *all* data of that particular type. For instance, if someone requested data about, say, GPs' earnings in Welwyn Garden City, the data to be considered for release should be about all GPs' earnings.

In other words, an information entrepreneur should not be able to frame a request to him or her a competitive advantage over others via an information asymmetry. If the request was granted, the result should be an increase in *public* good.

Recommendation 6: Create sector transparency panels

Requests for data should be considered by a competent body. This should be below the level of the Transparency Board, which should keep its strategic advisory role.

The most logical step is to create dedicated transparency panels distributed across sectors. These bodies should determine, among other things, whether there was a *prima facie* privacy threat from the release. If not, then there would be no privacy objection to the information's release, and an instruction should be sent to that effect to the government department or agency controlling the data, which should then work to place the dataset on its release schedule. The panels could help manage demand by influencing release schedule ordering.

In most cases, it would be helpful for a sector panel to cross ministerial or agency boundaries, to prevent the panel being 'captured' by a particular ministry or agency. However in some areas (e.g. transport), that may not be possible.

I will not make any specific recommendation about the size or composition of sector panels, or indeed which sectors should be served with panels. It is worth experimenting to determine best practice. Furthermore, it may be that different sectors have different requirements (e.g. in health, it may be best to proceed cautiously and rigorously, and so to have a sector panel with a wide and diverse membership, whereas in other areas, e.g. transport, the panel could be smaller and nimbler).

However, the procedure I shall suggest in recommendation 7 will influence the composition of the panels, or at least the rosters of experts upon whom they will draw. For instance, each panel should have access to technical advisors if it did not already include them in its membership.

Recommendation 7: A procedure for pre-release screening of data to ensure respect for privacy

For reasons set out in section 1.3.1, I shall not be recommending any strict set of rules or institutions. Rather, I set out a broad outline of a procedure that could be implemented in a number of ways. The aim is to suggest a method of pre-release screening that will work in a variety of contexts, including currently, as well as in a more demand-driven regime.

Note that this means that I will not provide a method of squaring the circle of releasing deanonymisable data on the Web. This is not a problem that can be *solved*; it involves a set of risks and potential benefits that can be *evaluated*. This evaluation will depend on the current state of knowledge, public opinion and political preferences. It will depend on whether the transparency programme has proved a success or not. It will depend on public attitudes to privacy, which are evolving very rapidly all the time.

The debate and discussion should be on a case-by-case basis, and hence this recommendation will not include criteria for making the decision, or red lines which should never be crossed. These will be a matter for policymakers and public at the time of a data release.

Note also that there is a nascent infrastructure in place already that might be built upon – OPSI's public sector information unlocking service (<http://unlockingservice.data.gov.uk/>) is designed to allow people to ask for information, and also has a screening service based on the Freedom of Information Act principles.

How would such a procedure work? I envisage a sequence of stages such as the following. To repeat, the exact set of institutions that implements these procedures need not be specified here, especially as the transparency programme is in an early phase and is bound to evolve over time. And although I have assumed a demand-driven transparency process, this procedure, beginning at stage C, would also be appropriate for the current top-down context.

A. Maintenance of a data asset register

See recommendation 5.

B. User demand

Information entrepreneurs should be able to make requests for data from the data asset register, on the bases established in recommendation 5. At this stage, the entrepreneur need not disclose the purpose of the request.

C. Screening of user requests

A competent body, which could be the sector panel as envisaged in recommendation 6, or alternatively could be an *ad hoc* body accredited by the panel, should then screen the request.

In the event that the body finds a *prima facie* threat, then the sector panel should be able to convene an inquiry to consider the data.

D. Consideration of potential privacy threat

If there was a potential threat, this should be assessed by the inquiry. The assessment should consist of rigorous analysis and debate with *all relevant stakeholders*. These

stakeholders consist of representatives of at least the following five groups of interested parties.

- a. Those tasked with increasing transparency and publishing public data.
- b. Those tasked with protecting privacy.
- c. Domain experts.
- d. Technical experts with understanding of deanonymisation techniques and the current threat model.
- e. The information entrepreneur(s) who made the original request.

The role of the domain expert is to provide expertise about how data are, or could be, used in the specific sectoral context – e.g. health, education, university education, primary education, etc. The role of the information entrepreneur is to explain the role that he or she envisages the data will play, and to demonstrate how any privacy concerns will be allayed.

The particular roles could be allocated to specific individuals ahead of time, or a group could be assembled *ad hoc*.

The value of debate about this issue was discussed in section 3.2.3, and the content in section 3.2.2. The structure of debate would be a matter for the chairperson. However, the heuristics provided by Nissenbaum (2010, and see section 3.1.4) give a sense of what questions would be relevant and how they should be framed. The issues should look at norms and expectations in the domain, and in that context, consider the risks and benefits of a data release.

The debate should as far as possible combine rigour with convenience and cheapness. It could take the form of a face-to-face meeting around a table, an extended email exchange, or collaboration on a written report.

One possible output would be a privacy impact assessment (PIA). Particularly in a demand-driven transparency regime, there is an issue as to who should accept the costs of a PIA associated for a particular release of data. Should it be the data provider who will be accountable for a reckless release? Or the entrepreneur demanding the data? If the latter, then that will clearly increase the costs of asking for data, as any request may of course be turned down. Furthermore, the entrepreneur's incentives would all be for transparency and against privacy if there was a clash, possibly compromising the legitimacy of any PIA he or she commissions.

Hence one possible structure is that, if a potentially privacy-threatening release of data is planned, the relevant stakeholders listed above should meet and together produce a PIA for consideration by data managers.

There is no reason to think that this will be a costly procedure in the large. Best practice and precedents will become available, allowing shortcuts in future decision-making and clarifying amendments to the data asset register. At present it is of course unknown what proportion of requests would need to be dealt with by an inquiry, and what proportion of those would be hard cases that would require hours of deliberation.

It should also be noted that the way that the current transparency regime is working, the number of borderline cases is actually very small, both in number and as a proportion of scheduled releases. Very few datasets containing data that could become

personal data are scheduled for release. Therefore, on current assumptions, this assessment procedure is unlikely to be triggered very often.

If it turned out that this procedure *was* prohibitively expensive, then that could be dealt with extremely easily by raising the bar and turning down more requests at an earlier stage. If privacy questions were fundamentally difficult and expensive to resolve, then the transparency programme should err on the side of caution and take fewer risks with privacy. This would mean fewer data releases, but also less risk to privacy.

E. Decision

The inquiry should either (i) reach a decision about whether, and if so how, the data should be released, or (ii) provide a report upon which a decision should be based. The decision should include specification of any level of aggregation, pseudonymisation or anonymisation required. It may be that aggregated data become useless for the entrepreneur's purposes, because the low-level patterns that he or she is looking for will disappear in the aggregation process. Hence it is vital that the entrepreneur is represented in the discussion.

It should also include a recommendation about whether any access or disclosure control system (cf. section 4.3.2 and recommendation 2) should be imposed, and if so what that system would contribute. As noted in recommendation 2, access control systems will act in the interest of transparency if they enable data to be released that could not otherwise be released for privacy or other reasons.

The decision, and its grounds, should be made absolutely clear and should be transparent and open to scrutiny. The likely context of use should be stated, from information supplied by the entrepreneur. If it has been deemed necessary to invoke one of the exceptions listed in Article 8, then the ground for this invocation should be made clear. The technical underpinning of the decision should also be evident.

A sufficient period should be made available for public and media scrutiny of the decision before the data are actually released.

A procedure of the type envisaged here would have a number of important advantages.

1. It would ensure that privacy was protected, that the public interest was properly considered, and that an auditable debate trail would exist in the event that a legal process was eventually triggered. The debate trail would also be an important resource for media scrutiny, and in the preservation of public confidence in the transparency programme (cf. section 3.2.3).
2. Each case would be considered on its own merits. Case-by-case analysis is far more potent than overarching principles which are hard enough to create and which, if the arguments of Ohm and others are persuasive, will never be sensitive to individual circumstances (cf. section 1.3.2).
3. The arguments will be addressed, under this procedure, in advance of any harms that may result, resulting (one hopes) in prevention rather than identification of fault after the fact.
4. The inclusion of technical experts ensures that their understanding of the threats and opportunities in this space is not neglected (cf. recommendation 3).

Recommendation 8: Extend the research base and maintain an accurate threat model

As noted many times in this review, our understanding of the threats from deanonymisation, both theoretically and empirically, is relatively low. Furthermore, as a new area of government activity, our understanding of the politics and public perceptions of transparency is also slim-to-vanishing.

It is therefore likely to be worthwhile convening a working group containing practitioners from industry to help build up the evidence base about cutting-edge techniques for risk management in this area, and also about the threats to privacy from the latest data analysis techniques. This working group could be managed jointly with a professional or learned society.

More traditional academic research would be very valuable in a number of areas. One area where the United Kingdom lags, particularly relative to the United States, is the mathematics and computer science of privacy. There are a number of world-leading experts in the field based in the UK, but cutting-edge work on topics such as differential privacy and deanonymisation tends to originate in the US. Privacy science is a burgeoning field which is clearly of great importance, and experts in the field will become extremely valuable and influential over time. It should become a research priority.

It is important to boost research into privacy in the context of open systems such as the World Wide Web, rather than in controlled systems which allow disclosure, query and access control. The mathematics of the latter are better-behaved, but open systems are the future. The research programme of Web Science has been in place for a small number of years, to develop theories of the Web as an unbounded sociotechnical system, and, from the point of view of the issues raised in this report, privacy science is quite naturally understood as a branch of Web Science.

One area where the United Kingdom is in advance of many other nations is research into practicable methods for consent management. This is a very hard problem, but ongoing projects such as ENCORE and VOME are developing a multi-disciplinary foundation for research in this area. Progress in understanding and managing consent will enable corresponding progress in several areas of transparency, not restricted to transparency programmes of the type discussed in this report. For instance, there are a number of important issues with respect to the sharing of personal data with controls to third parties, where effective consent management could help preserve privacy while increasing the quantity and quality of legitimate data sharing.

There are also some low-cost methods of extending crucial areas of understanding, exploiting already-existing mechanisms.

1. **Datalabs.** In order to help assess how easy/difficult it is to deanonymise a proposed data release, datalabs (such as are already run by a number of government agencies, including the Office for National Statistics and Her Majesty's Revenue and Customs, as well as the Secure Data Service of the ESRC-funded UK Data Archive) can be exploited. In a secure environment, researchers could be invited to deanonymise a dataset proposed for release, bringing whatever auxiliary information they liked to support their quest. In an extension of this method, academic researchers could be invited to do the same in competition, perhaps for a low-value prize – this sort of 'hackfest' is a well-understood institution in academe, where the main gain is academic kudos. I

should emphasise that this sort of exercise need not be performed for every dataset, but rather should be performed on real or realistic dummy datasets on a regular basis in order to educate policymakers on the extent of the real-world threat to privacy from jigsaw identification. The threat model currently in place lacks verification and is based to a large extent on intuition (for a rare example of the use of this technique to inform us empirically, see Lafky 2009). The use of secure datalabs and competitions will inform the threat model.

2. **Citizens' panels.** Many government departments and agencies conduct citizens' panels on a regular basis, testing public opinion on a number of matters. Even given the demand-driven regime of recommendation 4, representation of citizens' opinions on particular data releases would be somewhat second-hand, via the efforts of information entrepreneurs. Citizens' confidence in transparency is vital – consulting citizens' panels could help to gauge and track opinion on this somewhat complex issue, and also of disseminating awareness through the wider population.

Finally, as part of the threat model, we also need to understand what the business model for an adversary would be, and what harms might be precipitated. Whether an adversary could actually monetise the effort of deanonymising data is an important factor in assessing risk; at the moment, there seem to be few convincing examples of how someone may profit from discovering and revealing identities.

It is important, given the need for public confidence, that the transparency programme does not outpace our understanding of the threat, or our understanding of public attitudes. However, the need to be anchored in empirical understanding does not mean that the transparency programme should be risk averse. The deanonymisation risk is unlikely to be large, and public attitudes may well be less concerned with privacy than is often thought. Research to fill these gaps in our understanding could be done relatively quickly. However, in general, an incremental approach to data releases, and to levels of anonymisation within particular datasets, will be more privacy- and confidence-preserving in the absence of telling empirical results.

Recommendation 9: Create a guidance product to disseminate best practice and current research in transparency

The practice of transparency programmes is, as noted, very new. There will inevitably be a lot of trial and error, and reinvention of the wheel. Furthermore, the legal background is extremely sparse, with relatively little guidance provided by current data protection jurisprudence. A particular danger is that governments become too averse to risk. While they should be very much aware of the risks to, and properly protective towards, citizens' privacy, this should not lead them to retrenchment and failure to act. A number of experiments in transparency (of all kinds) have been in train across the world, and best practice will eventually emerge concerning problems practical, legal and moral (including ways of maximising use of data without compromising citizens' privacy).

Some of the recommendations in this report – in particular recommendations 7 and 8 – will be important contributions to best practice in this area. Experience in other nations, and other types of administration (local governments, city governments, health service providers, even commercial companies) will also contribute to our understanding of this burgeoning field.

The United Kingdom government could play an important part by aggregating this experience in a guidance product, highlighting:

- How authorities and those charged with data sharing manage the processes.
- How the law is evolving internationally.
- How threat models are evolving.
- Where economic gains and losses are to be made.
- How privacy is being preserved.

The product will be in effect a journal of best transparency practice. One of its obvious functions would be to publish, and to act as a repository of, the debate trails about privacy and transparency that would occur under recommendation 7. It would also be an obvious route to publication for the research that I advocate in recommendation 8. The specific issues of consent management and privacy science, discussed in recommendation 8, should be included in the remit.

Such a guidance product would be a boon in an area where activity is somewhat diffuse, and would help create coherence and common standards. The development of standards could prompt nations or organisations with fewer resources to become more transparent. The product would also be an important resource for benchmarking progress in the UK.

The recommendation does not dictate the form of the product. One possibility would be a website, another would be a magazine existing both online and offline. It could either attract original contributions, or aggregate existing ones (or both). The contributions should not be technical, though they should be technically-informed; the product should communicate best practice to practitioners. The product should be open access.

A high-profile and respected editor would be important, and at least the major original contributions should be peer-reviewed. A high quality editorial board, made up not

only of academics, but also key industrial partners, should manage the product. It may be that a learned society could be included as a partner for management and delivery.

Recommendation 10: Keep the efficacy of control in the new paradigm under review

Under the current retentive Whitehall culture, privacy receives a lot of protection via the practical obscurity engendered by the sheer difficulty of getting information released in a scheduled manner. Privacy is protected at least partly because civil servants are reluctant to part with data. Rules for information control go with the grain of this retentive culture, and so can be expected to be relatively effective. Even so, there have been plenty of well-publicised errors with laptops left on trains.

If the transparency programme achieves its ultimate aims, then there will be a culture shift in Whitehall towards a paradigm of transparency and sharing. The default question to be asked will no longer be 'why should I release these data?' but rather, 'why should I *not* release these data?' In such a culture, privacy-preserving rules for information control will go against, rather than with, the grain. It may be that the current rules, adequate in a retentive culture, will no longer be adequate in a transparent culture. Civil servants tend to err on the side of caution. Under the current paradigm, caution = retention. Under a new transparent paradigm, caution = publication.

Under any regime, one of the greatest threats to privacy where data are collected is the accidental (or even malicious) release of data outside the rules, and this will always be a risk however carefully crafted the rules are. It should never be assumed that rules will be followed perfectly on every occasion.

Departments and agencies should ensure that the rules for retaining information whose release would threaten privacy remain adequate in a world in which the natural assumption is that information should be released. The situation should be kept under review.

Recommendation 11: Maintain existing procedures for identifying harms and remedies

The question of establishing harms and remedies needs to be considered along the lines argued in section 3.3. There are two avenues open to the aggrieved citizen at the moment. First, he or she can apply to the courts for legal redress. This is not a satisfactory situation, because in the area of privacy, this provides easier access for a rich person than someone of average income.

The alternative is to apply to the ICO for help, advice and redress. This is clearly a more accessible route than the courts, and should stay in place.

However, it may be that if a large number of cases is generated by the transparency programme, the ICO's resources will be strained. In that case – and it must be emphasised that there is currently no evidence that this will be the case – then additional means for assessing harms and determining remedies must be sought.

The ICO does have leeway to prioritise particular cases, and to treat related complaints as a group, so that a series of similar cases (people affected by the same release of data) is judged together. It therefore has some means of coping with a rapid increase in complaints.

The ICO does, however, only have powers to fine, not to award compensation. This again should not be immediately considered as problematic, as privacy breaches tend to cause loss of reputation or station, rather than financial losses. These may be irreparable, but may not be remediable by financial compensation. However, until we have seen the effects of clashes between privacy and transparency, we cannot be sure that this will be the case in future. For instance, the determination of the government to publish details of contracts with private-sector suppliers might result in financial losses that would not otherwise have occurred. In such cases the ICO, though it would be able to criticise and fine departments or agencies for malpractice, would not be able to award compensation – the complainant would have to seek redress through the courts.

If the ICO's workload expanded unreasonably, or if it became clear that many of those whose privacy was breached had reasonable claims for compensation, then it may be that a new institution should be created to deal with complaints generated by the transparency programme. It might be worth considering creating a data ombudsman to adjudicate cases and award compensation where deserved.

However, until it is shown that the transparency programme does generate a number of complaints which cannot be effectively served under the current regime, no action should be taken. Note also that if its technical expertise is boosted as recommended in recommendation 3, the ICO would be better able to determine the recklessness or otherwise of a release of a dataset on the Web.

The situation should be kept under review, but at present the current regime should remain in place.

Recommendation 12: Use data.gov.uk to raise awareness of data protection responsibilities

Finally, the Data Protection Act applies to people downloading data from data.gov.uk. Whether downloading data counts as an act of processing under the Act is perhaps doubtful (is it an act of 'retrieval'?). If the data are not used by entrepreneurs in any useful way, or only for domestic purposes, or are non-personal data then presumably no duties are incurred by them. However, if someone does manipulate personal data and disseminate the resulting output, then they become a data controller in the meaning of the Act, and incur the relevant duties. In particular, if they take aggregated or anonymised data that have been released and reidentify individuals from that data, they clearly fall under the aegis of the Act.

data.gov.uk should include prominent and clear reminders to those downloading datasets of the provisions of the Data Protection Act, and clearly state that best practice includes not attempting to deanonymise data. The reminders should be prominent enough to register, should include links to the Act, and yet should not interfere with or prevent the downloading of the data. For example, they could take the form of banner advertisements.

Of course this final suggestion will not deter someone who is determined to deanonymise data for anti-social reasons, and will have little effect on someone outside EU jurisdiction or the OECD. But it has a relatively low cost, may prevent inadvertent breaches of privacy, and will help spread awareness of data protection law.

Recommendation 13: Investigate the Vulnerability of Anonymised Databases

To accept the arguments of (Ohm 2010) in full would be to neutralise most of the positive benefits of the transparency programme. However, the potential to deanonymise, and the impossibility of exercising access or query control, mean that releasing anonymised datasets will constitute a small risk to privacy. The extent of this risk is unclear; although it is currently extremely low, it may not be negligible. If the empirical work to develop the threat model which I recommend above is carried out, then our understanding of the risk will become clearer.

Given the lack of certainty, and given the transparency programme's current focus on extracting value from anonymised datasets, it would be very valuable to scrutinise the latest sophisticated anonymisation techniques, alongside a critique of the practical capabilities of known deanonymisation methods. A working group with the requisite technical abilities should be convened to investigate the extent of the risk along the lines suggested in recommendation 8. On the basis of current knowledge (cf. Lakfy 2009), the group's task will be to confirm the low level of risk, and therefore to provide extra reassurance to those with a legitimate privacy concern. If the risk appears to be larger than the evidence currently suggests, then a rethink will of course be required.

One cheap, quick and direct method of testing the risk would be to commission researchers to try to identify individuals from sample datasets under secure conditions. Whether or not the datasets could be 'cracked', the result of such an exercise would be a greater understanding of the threat – an undeniable benefit to both privacy and transparency campaigners. In particular, valuable information could be gathered about where vulnerabilities exist, which parameters are most helpful to those trying to deanonymise data, and which levels of aggregation are most appropriate for preserving the anonymity of data subjects.

As noted in recommendation 8, transparency practice should not outpace our understanding of the risk. The threat is almost certainly very low, but it is important that that intuition is backed up with evidence.

Recommendation 14: Be transparent about the use of anonymisation techniques

Given the importance of transparency about the transparency programme (section 3.4), and the value of debate and research about the deanonymisation threat (recommendations 7, 8 and 13), it will be extremely helpful if agencies are open about the techniques that they use to anonymise, pseudonymise or perturb datasets. This will facilitate sensible and accurate debate about the risks and benefits of data releases (and will not, of course, enable an adversary to reverse-engineer the original database of personal data).

6 Acknowledgements

The reviewer's thanks go to the following people and organisations who engaged with and advised the review process. Many experts will recognise their ideas included in the text above, and I am grateful to everyone who gave their time to see me, and who invited me to events from which I invariably learned a great deal. Of course, the responsibility for the use that I have made of information and ideas lies solely with me, and the appearance of a person or organisation in these acknowledgements should not be taken to imply their agreement with the arguments of this report.

Representatives of the following public departments, agencies and offices

Association of Chief Police Officers

Cabinet Office

Department for Education

Equality and Human Rights Commission

Her Majesty's Revenue and Customs

Home Office

Information Commissioner's Office

Local Public Data Panel

Ministry of Justice

The National Health Service Information Centre for Health and Social Care

National Police Improvement Agency

No.10 Downing Street

Office for National Statistics

Transparency Board

Experts on privacy and/or transparency

Ross Anderson, Cambridge University

Paul Aylin, Imperial College, Dr Foster Unit

Caspar Bowden, Independent Consultant (previously Microsoft)

Ian Brown, Oxford Internet Institute

Conn Crawford, Sunderland City Council

Neil McBride, De Montfort University

Helen Nissenbaum, New York University

Will Perrin, Talk About Local

Chris Pounder, Amberhawk

Charles Raab, University of Edinburgh

Edgar Whitley, London School of Economics

Representatives of the following non-governmental organisations

Centre for Policy Studies

Demos

IMS Health

Liberty

NO2ID

Open Rights Group

Privacy International

Victim Support

Which?

Participants in the following events

VOME workshop on Delivering Public Services Online, Dec 7th, 2010

Demos workshop on The Promise and Perils of Digital Era Governance, Jan 26th, 2011

Guardian workshop on Managing Public Sector Information, Mar 1st, 2011

Workshop on Open Government Data and Privacy, Catholic University of Leuven, Mar 9th, 2011

Intellect workshop on transparency and privacy, Mar 15th, 2011

Information Commissioner's Office workshop on Privacy and Data Anonymisation, Mar 30th, 2011

Administration Data Liaison Service's workshop on Showcasing UK Administrative Data: Innovative Research, Methodologies and Modes of Access, Royal Statistical Society, Jul 11th, 2011

I should like to thank my employers at the School of Electronics and Computer Science at the University of Southampton, for allowing me the time and space to conduct this review. Thanks also to Derine Nelson-Streeter of the Cabinet Office for invaluable administrative and diary support,.

7 References

Alastair Allan & Pete Warden (2011). 'Got an iPhone or 3G iPad? Apple is recording your moves', *O'Reilly Radar*, <http://radar.oreilly.com/2011/04/apple-location-tracking.html>.

Ross Anderson (2008). *Security Engineering: A Guide to Building Dependable Distributed Systems*, 2nd edition, Indianapolis: Wiley.

Article 29 Data Protection Working Party (2007). *Opinion 4/2007 on Concept of Personal Data*, http://ec.europa.eu/justice/policies/privacy/docs/wpdocs/2007/wp136_en.pdf.

Jane Bailey & Ian Kerr (2007). 'Seizing control? The experience capture experiments of Ringley & Mann', *Ethics and Information Technology*, 9, 129-139.

BBC Online (2010). 'German vandals target Street View opt-out homes', 24th Nov, 2010, <http://www.bbc.co.uk/news/technology-11827862>.

Helen Beetham, Lou McGill & Allison Littlejohn (2009). *Thriving in the 21st Century: Learning Literacies for the Digital Age (LLiDA Project)*, <http://www.jisc.ac.uk/media/documents/projects/llidareportjune2009.pdf>.

Christian J. Bonnici & Lizzie Coles-Kemp (2010). 'Principled electronic consent management: a preliminary research framework', *Proceedings of the 2010 International Conference on Emerging Security Technologies*, <http://doi.ieeecomputersociety.org/10.1109/EST.2010.21>.

danah boyd (2008). *Taken Out of Context: American Teen Sociality in Networked Publics*, PhD thesis, University of California, Berkeley.

Peter Bradwell (2010). *Private Lives: A People's Inquiry into Personal Information*, London: Demos.

Ian Brown (2010). 'Data protection: the new technical and political environment', *Computers and Law*, 20.

Ian Brown, Lindsey Brown & Douwe Korff (2010). 'Using NHS patient data for research without consent', *Law, Innovation and Technology*, 2(2), 219-258.

Cabinet Office (2010). *Coalition Agreement*, http://www.cabinetoffice.gov.uk/sites/default/files/resources/coalition_programme_for_government.pdf.

Cabinet Office (2011a). *Transparency: Publication of New Central Government Contracts*, guidance note, Feb 2011, <http://www.cabinetoffice.gov.uk/sites/default/files/resources/Guidance%20-%20Publication%20of%20new%20central%20government%20contracts%20-%20February%202011.pdf>.

Cabinet Office (2011b). *Transparency: Publication of Tender Documentation*, guidance note, Feb 2011, <http://www.cabinetoffice.gov.uk/sites/default/files/resources/Guidance%20-%20Publication%20of%20new%20central%20government%20tender%20documents%20-%20February%202011.pdf>.

David Cameron (2010). *Letter to Government Departments on Opening Up Data*, <http://www.number10.gov.uk/news/statements-and-articles/2010/05/letter-to-government-departments-on-opening-up-data-51204>.

David Cameron (2011). *Letter to Cabinet Ministers on Transparency and Open Data*, <http://www.number10.gov.uk/news/statements-and-articles/2011/07/letter-to-cabinet-ministers-on-transparency-and-open-data-65383>.

Ann Cavoukian & Khaled El Emam (2011). *Dispelling the Myths Surrounding De-identification: Anonymization Remains a Strong Tool for Protecting Privacy*, Ontario: Office of the Privacy and Information Commissioner, <http://www.ipc.on.ca/images/Resources/anonymization.pdf>.

Lizzie Coles-Kemp & Elahe Kani-Zabihi (2010). 'On-line privacy and consent: a dialogue, not a monologue', *Proceedings of the 2010 Workshop on New Security Paradigms (NSPW'10)*, <http://dx.doi.org/10.1145/1900546.1900560>.

Lizzie Coles-Kemp, Yee-Lin Lai & Margaret Ford (2010). *Privacy on the Internet: Attitudes and Behaviours*, London: VOME.

Gareth Crossman with Hilary Kitchin, Rekha Kuna, Michael Skrein & Jago Russell (2007). *Overlooked: Surveillance and Personal Privacy in Modern Britain*, London: Liberty, <http://www.liberty-human-rights.org.uk/policy/reports/overlooked-privacy-report-december-2007.pdf>.

Tore Dalenius (1977). 'Towards a methodology for statistical disclosure control', *Statistik Tidskrift*, 15, 429-444.

Paul B. de Laat (2008). 'Online diaries: reflections on trust, privacy and exhibitionism', *Ethics and Information Technology*, 10, 57-69.

Denis Diderot ([1775]1995). 'Encyclopédie', in Isaac Kramnick (ed.), *The Portable Enlightenment Reader*, New York: Penguin, 17-21.

Cynthia Dwork (2006). 'Differential privacy', in *Proceedings of 3rd International Colloquium on Automata, Languages and Programming (ICALP)*, Berlin: Springer, 1-12.

Cynthia Dwork (2008). 'Differential privacy: a survey of results', in *Proceedings of the 5th International Conference on Theory and Applications of Models of Computation (TAMC)*, Berlin: Springer, 1-19.

Amitai Etzioni (1999). *The Limits of Privacy*, New York: Basic Books.

Archon Fung, Mary Graham & David Weil (2007). *Full Disclosure: The Perils and Promise of Transparency*, New York: Cambridge University Press.

Simson L. Garfinkel (2008). 'Information of the world, unite!' *Scientific American*, Sept 2008, 60-65.

Charles E. Gilbert (1959). 'The framework of administrative responsibility', *Journal of Politics*, 21, 373-407.

F.A. Hayek (1945). 'The use of knowledge in society', *American Economic Review*, 35(4), 519-530.

Information Commissioner's Office (2011). *Data Sharing Code of Practice*, http://www.ico.gov.uk/for_organisations/data_protection/topic_guides/~media/docu

[ments/library/Data Protection/Detailed specialist guides/data sharing code of practice.ashx](#).

Adam N. Joinson & Carina B. Paine (2007). 'Self-disclosure, privacy and the Internet', in Adam Joinson, Katelyn McKenna, Tom Postmes & Ulf-Dietrich Reips (eds.), *The Oxford Handbook of Internet Psychology*, Oxford: Oxford University Press, 237-252.

Maria Karyda & Spyros Kokolakis (2008). 'Privacy perceptions among members of online communities', in Alessandro Acquisti, Stefanos Gritzalis, Costas Lambrinouidakis & Sabrina De Capitani di Vimercati (eds.), *Digital Privacy: Theory, Technologies and Practices*, Boca Raton, FL: Auerbach Publications, 253-266.

Tim Kelsey (2009). 'Long live the database state', *Prospect*, 161.

Kenneth Kernaghan (2000). 'The post-bureaucratic organization and public service values', *International Review of Administrative Sciences*, 66, 91-104.

Douwe Korff (2003). *Report on the Findings of the EC Study on the Implementation of the Data Protection Directive*, European Commission, 2003, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1287667.

Deborah Lafky (2009). *The Safe Harbor Method of De-Identification: An Empirical Test*, Office of the National Coordinator for Health Information Technology, http://www.ehcca.com/presentations/HIPAAWest4/lafky_2.pdf.

Ninghui Li, Wahbeh Qardaji & Dong Su (2011). 'Provably private data anonymization: or, k-anonymity meets data privacy', <http://arxiv.org/abs/1101.2604v1>.

Steve Lohr (2010). 'Netflix cancels contest after concerns are raised about privacy', *New York Times*, 12th Mar, 2010.

Francis Maude (2010). 'Our plans for government transparency continue apace – but we must uphold the highest possible standards of personal data protection', conservativehome.com, 15th Oct, 2010, <http://conservativehome.blogs.com/platform/2010/10/francis-maude-mp-our-plans-for-government-transparency-continue-apace-but-we-must-uphold-the-highest.html>.

Derek McAuley, Hanif Rahemtulla, James Goulding & Catherine Souch (2011). 'How open data, data literacy and linked data will revolutionise higher education', in Louis Coiffait (ed.), *Blue Skies: New Thinking About the Future of Higher Education*, London: Pearson, 88-93, <http://pearsonblueskies.com/how-open-data-data-literacy-and-linked-data-will-revolutionise-higher-education/>.

Frank McSherry (2010). 'Privacy integrated queries: an extensible platform for privacy-preserving data analysis', *Communications of the ACM*, 53(9), 89-97.

John Stuart Mill (1859). *On Liberty*, London: John W. Parker & Son.

John Stuart Mill (1861). *Considerations on Representative Government*, London: Parker, Son & Bourn.

Arvind Narayanan (2009). 'De-anonymization is not X: the need for re-identification science', *33 Bits of Entropy* blog, 14th Oct, 2009, <http://33bits.org/2009/10/14/de-anonymization-is-not-x-the-need-for-re-identification-science/>.

Arvind Narayanan & Vitaly Shmatikov (2008). 'Robust de-anonymisation of large sparse datasets', in *Proceedings of the 2008 IEEE Symposium on Security and Privacy* 111-125.

Arvind Narayanan & Vitaly Shmatikov (2009). 'De-anonymizing social networks' in *Proceedings of the 2009 IEEE Symposium on Security and Privacy*, 173-187.

Arvind Narayanan & Vitaly Shmatikov (2010). 'Myths and fallacies of "personally identifying information"', *Communications of the ACM*, 53(6), 24-26.

Helen Nissenbaum (2010). *Privacy in Context: Technology, Policy, and the Integrity of Social Life*, Stanford, CA: Stanford University Press.

OECD (1980). *OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data*, http://www.oecd.org/document/18/0,3343,en_2649_34255_1815186_1_1_1_1,00.html.

Kieron O'Hara (2010). 'Intimacy 2.0: privacy rights and privacy responsibilities on the World Wide Web', in *Proceedings of the Web Science Conference 2010*, Raleigh NC, <http://eprints.ecs.soton.ac.uk/18760/>.

Paul Ohm (2010). 'Broken promises of privacy: responding to the surprising failure of anonymization', *UCLA Law Review*, 57, 1701-1777.

Gavin Phillipson (2006). 'The "right" of privacy in England and Strasbourg compared', in Andrew T. Kenyon & Megan Richardson (eds.), *New Dimensions in Privacy Law: International and Comparative Perspectives*, Cambridge: Cambridge University Press, 184-228.

William L. Prosser (1960). 'Privacy', *California Law Review*, 48, 338-423.

Charles D. Raab (1999). 'From balancing to steering: new directions for data protection', in Colin J. Bennett & Rebecca Grant (eds.), *Visions of Privacy: Policy Choices for the Digital Age*, Toronto: University of Toronto Press, 68-93.

Charles D. Raab (2005). 'The future of privacy protection', in Robin Mansell & Brian S. Collins (eds.), *Trust and Crime in Information Societies*, Cheltenham: Edward Elgar, 282-318.

Charles D. Raab (forthcoming). 'Privacy, social values and the public interest', *Politische Vierteljahresschrift*.

Beate Rössler (2005). *The Value of Privacy*, Cambridge: Policy Press.

M. Angela Sasse & Ivan Flechais (2005). 'Usable security: why do we need it? How do we get it?' in Lorrie Faith Cranor & Simson Garfinkel (eds.) *Security and Usability: Designing Secure Systems That People Can Use*, Sebastopol CA: O'Reilly Media, 13-30.

Daniel J. Solove (2008). *Understanding Privacy*, Cambridge, MA: Harvard University Press.

John T. Soma & Stephen D. Rynerson (2008). *Privacy Law in a Nutshell*, St Paul, MN: Thomson/West.

Latanya Sweeney (1997). 'Weaving technology and policy together to maintain confidentiality', *Journal of Law, Medicine and Ethics*, 25, 98-110.

Latanya Sweeney (2002). 'k-anonymity: a model for protecting privacy', *International Journal on Uncertainty, Fuzziness and Knowledge-Based Systems*, 10, 557-570.

Latanya Sweeney (n.d.). *Computational Disclosure Control: A Primer on Data Privacy Protection*, <http://groups.csail.mit.edu/mac/classes/6.805/articles/privacy/sweeney-thesis-draft.pdf>.

Transparency Board (2010). *Public Data Principles*, http://data.gov.uk/wiki/Public_Data_Principles.

Raymond Wacks (2006). 'Why there will never be an English common law privacy tort', in Andrew T. Kenyon & Megan Richardson (eds.), *New Dimensions in Privacy Law: International and Comparative Perspectives*, Cambridge: Cambridge University Press, 154-183.

Raymond Wacks (2010). *Privacy: A Very Short Introduction*, Oxford: Oxford University Press.

Ian Walden (2007). 'Privacy and data protection', in Chris Reed & John Angel (eds.), *Computer Law: The Law and Regulation of Information Technology*, 6th ed., Oxford: Oxford University Press, 459-504.

Michael White (2011). 'Crime maps: too much information?' *Guardian Politics Blog*, 1st Feb, 2011, <http://www.guardian.co.uk/politics/blog/2011/feb/01/crime-maps-too-much-information?intcmp=239>.

Edgar A. Whitley (2009). 'Informational privacy, consent and the "control" of personal data', *Information Security Technical Report*, 14(3), 154-159.

Jenny Williams (2010). 'EC delays revision of data protection directive', *Computer Weekly*, 5th Aug, 2010, <http://www.computerweekly.com/Articles/2010/08/05/242267/EC-delays-revision-of-data-protection-directive.htm>.

James Q. Wilson (1980). *The Politics of Regulation*, New York: Basic Books.