



May 2016

Public dialogue on the ethics of data science in government

Government
Data Science
Partnership

Contents

1 Executive summary	2
2 Introduction and methodology	8
2.1 Background	8
2.2 Objectives	9
2.3 Methodology	9
2.4 Interpretation of data	15
2.5 This report	17
2.6 Acknowledgements	17
3 Data science in context.....	19
3.1 Generating data	19
3.2 From data to data science	20
3.3 Perceptions of government use of data science	23
3.4 Towards opportunities for data science	26
4 Engaging the public in data science.....	29
4.1 Importance of case studies when discussing data science	29
4.2 Quantitative experiments	30
4.3 Future engagement.....	33
5 Priorities and opportunities	37
5.1 Assessing overall value of data science project.....	39
5.2 Assessing where benefits outweigh risks	46
5.3 Exploring the features of data science projects	55
5.4 Implications for data scientists and policymakers	63
6 Differing perspectives.....	66
6.1 Public perspectives	67
6.2 Does experience of sharing data breed favourability?.....	69
7 Ethical framework and governance.....	73
7.1 Developing ethical principles	73
7.2 Transparency	74
7.3 Standards and mechanics	75
7.4 Data security	78
7.5 Outcomes and benefits	79
7.6 Recommendations for the Ethical Framework	80

1. Executive summary

1 Executive summary

1.1 Introduction

Over the last year, the Government Data Science Partnership¹ has been developing a Framework for Data Science Ethics.² The framework seeks to help policy makers and data scientists maximise the potential for data science within government, whilst navigating the legal and ethical issues that accompany new forms of data analysis.

Ipsos MORI was commissioned to undertake a public dialogue to inform further development of the ethical framework. With oversight from GDS, Sciencewise, GO-Science and ONS, the project sought to explore how the Government Data Science Partnership can maximise the opportunities of data analysis in an appropriate way.

The project consisted of three methods of public engagement:

- A series of **deliberative public workshops**³ to gather qualitative insight into public opinion on the appropriate use of data science. This included the development of a series of in-depth case studies and hypothetical abstracts to introduce the concept of data science and explore ethical issues.
- An **online survey**⁴ to develop robust quantitative evidence on what the public think makes government data science projects appropriate. This included a conjoint exercise to identify the underlying factors driving attitudes and decisions about government use of Data Science.
- Use of the results of the online survey to develop an **engagement tool** to interact with a wider audience in a public debate around data science.

An annex is published alongside this report. This includes topline data from the online survey and conjoint exercise, and copies of the recruitment and fieldwork materials.

1.2 Key findings: engagement and awareness

- **Awareness of data science was limited.** Although participants were comfortable with the concept of sharing personal information to access services, they were unaware of the volume and rate at which they generate data from numerous sources across their everyday lives. They also had little knowledge about how individual data records can be anonymised or aggregated to generate insights.
- **Across both the workshops and online survey, participants were broadly willing to support the need for government to find new ways of using data.** They recognised the potential uses and relevance that the large volumes of data we now generate and collect have for public services.

¹ The Government Data Science Partnership includes the Government Digital Service, Office for National Statistics and the Government Office for Science.

² Data Science Ethics, 8th December 2015, <https://data.blog.gov.uk/2015/12/08/data-science-ethics/>

³ Reconvened workshops were conducted with 88 people across London, Taunton, Sheffield and Wolverhampton.

⁴ An online survey of 2,003 people aged 16-75 in Great Britain was conducted between 24th February and 7th March 2016 using the Ipsos MORI Access Panel

- However, though the concept received broad support, participants were often unsure and cautious about the practicalities of how data is, or could be, collected and analysed. They were also sceptical of why data science was being considered in some circumstances, and how results might be used by government for policy-making and planning services. With very low level of baseline awareness of how data science works in any context, many initially struggled to see the value of using computer analytics (as opposed to more traditional methods) without further information about how these work in practice.
- Demonstrating the potential impact of data science through real life case studies was therefore crucial to engaging the public in a discussion about opportunities for data science. Sharing this information made participants more willing to assess and comment on the concepts and mechanics of different data science projects. Only at this point could they identify what was and was not appropriate. Yet, while further information made them more open to the concept in general, it did not necessarily affect individual assessments of benefit and risk, nor guarantee support for data science projects.
- The public had other assumptions about data and data science. Beyond limited awareness of 'data science', participants initially struggled to trust in the potential value of data science. This included: low awareness of how data sets are collated; doubt as to whether computers can make better decisions than humans; caution about techniques that cluster individuals or use of correlations between datasets that initially appear unrelated; and ambiguity about the level of control and automation that can or should be given to a computer.
- Where there was concern about government using data science, this was mainly focused on data security and the significance of decisions that government could take. Opportunities for government use of data science were therefore judged with an acute awareness of the potential risks to those referenced in the dataset, and to those receiving the intervention that follows.
- Publication of the guidance offered by government on data science was more likely to reassure than concern people. Many wanted to see a commitment to transparency, and welcomed discussion about a future framework. Based on the experience of conducting the research, and on the feedback of participants, future engagement with the public on data science should:
 - Account for low level awareness of data science, both of the data and techniques used, and of the detailed ethical trade-offs encountered in design;
 - Use case studies of previous data science projects (both successful and unsuccessful) to demonstrate the potential value, pitfalls and impact of these methods, and to unlock a more meaningful ethical discussion;
 - Account for the difficulty in separating attitudes to specific policies, data, government and data science techniques;
 - Give tangible examples of what any underlying principles included in the ethical framework mean in reality, for example using hypothetical scenarios that are familiar and easy for the public can relate to;

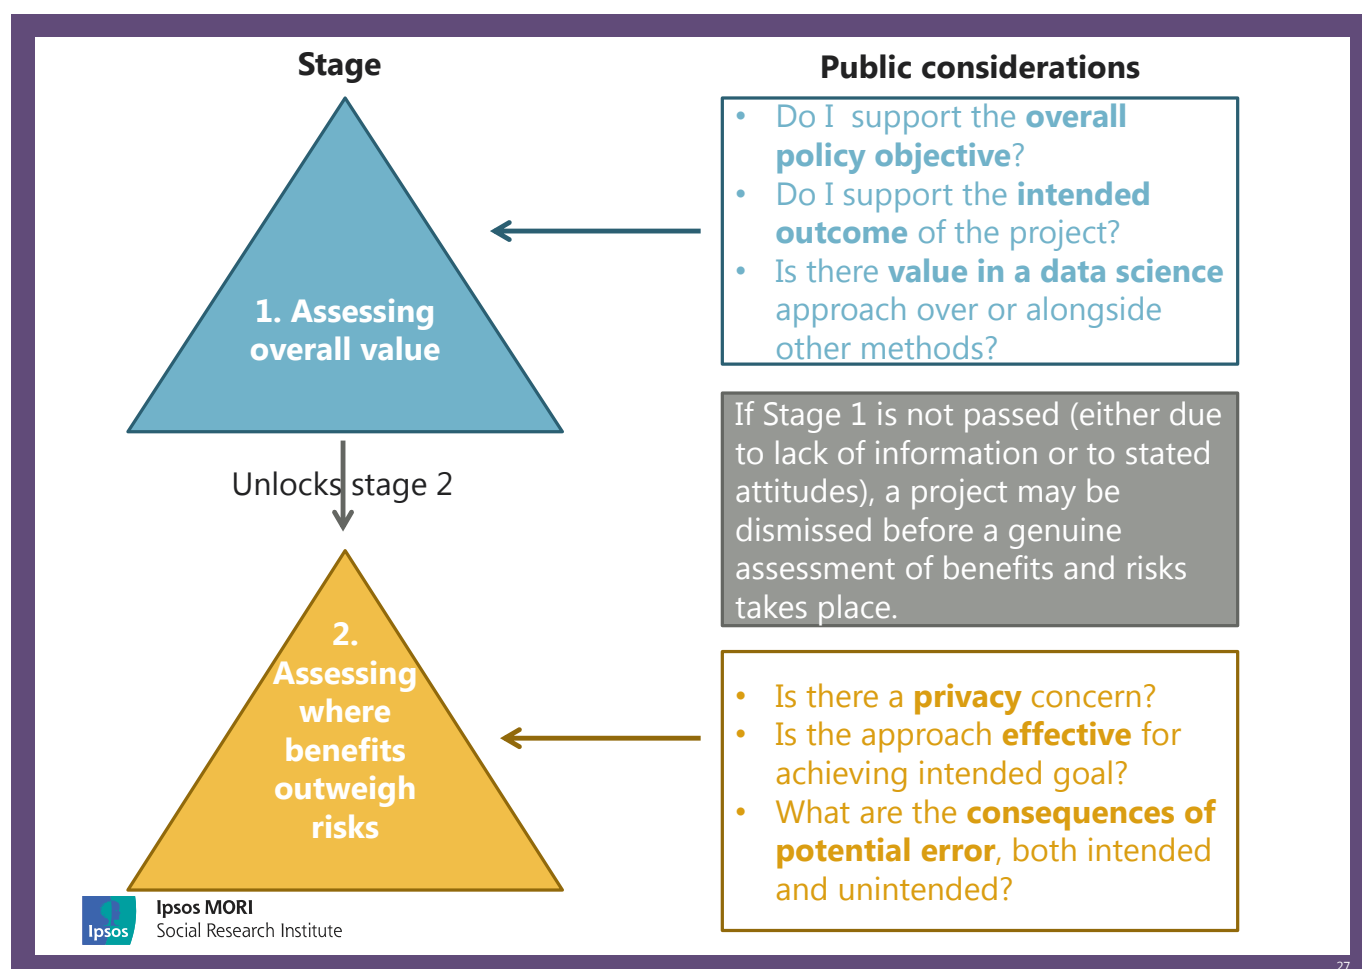
- Be aware that more education on data science opens up broad acceptability to government exploring opportunities for data science, but does not change the individual trade-off of risks and benefits applied for assessment of individual projects.

1.3 Key findings: opportunities for data science

- The public were broadly open to exploring opportunities for data science, though opinion was mixed. Close to half (47%) of adults from the survey were comfortable with government exploring new applications for data science; in contrast, just under a third (31%) were more cautious and thought that government should not explore ways of using data science due to privacy risks. Across both the survey and workshops, support for data science was conditional and varied substantially for individual data science scenarios.
- Opportunities for data science were seen as greatest where a clear wider public benefit could be established and where the risk of taking no action was significant. Before fully engaging in discussion about specific methodology, opportunities for data science needed to first pass a broader values assessment: Do I support the overall policy objective? Do I support the intended outcome of the project? Is there value in using a data science approach over or alongside other methods? When these conditions were not satisfied, or were contentious, many participants were inclined to dismiss projects outright before any real consideration of how data science could make an ethical and robust contribution.
- At this stage, the most obvious opportunities for data science emerged where there was a clear remit for government, where the current status quo was seen as inadequate, where the outcome was appropriate, and where data science could complement other methods. Early on in the discussion, opportunities for data science needed a clear benefit and value over any inconvenience or harm that may potentially occur as the result of error. This overall value assessment was made in the context of personal and ideological values about the role of government and value of data science. As set out in 1.2, the current understanding of the opportunities for data science was low and communications activity is required to address this.
- Once the public benefit and value of data science had been established, opportunities for data science were based on a nuanced risk assessment of the entire project. This assessment balanced three further considerations: Is there a privacy concern? Is the approach effective in achieving the intended policy goal? What are the consequences of potential error, either intended or unintended? Beyond assessment of public benefit, the next biggest driver in identifying opportunities for data science was *the type of data used*, followed by whether or not *individuals could be identified* within the data.
- There were no overall red-lines of acceptability; instead, concerns about risk, consequence and efficacy were measured against the specifics of the policy aim to judge the merits of a data science project. Here, small changes in the features of a data science project, within the context of a specific policy, could alter the overall assessment of acceptability. Some features had a higher baseline of acceptability than others, but all were considered in the context of the original policy objectives and intended outcomes. For example, using individual level data for more personalised careers services was more acceptable than for using individual level data to improve public transport.

- Low awareness of the technical aspects of data science meant that risks were often associated with the direct consequences of what happens when errors occur; however when prompted, participants were also able to judge different types of risk. Again, these were nuanced assessments based on: the nature of impact of the outcome for society and for the individual; the risk of doing nothing; and the number of people affected, and likelihood of errors occurring.

Figure 1.1: Assessing overall value before undertaking a nuanced risk assessment



1.4 Key findings: differing perspectives

- Perspectives on government use of data science differed noticeably between individuals. Future public engagement in exploring opportunities for data science should consider a number of factors and perspectives, in particular the impact of someone's lifestyle, past experience and political philosophy on how they receive data science in government. Differences did not align neatly to age; instead factors such as level of interaction with government services and ideological preferences about the relationship between citizen and state carried more weight.
- An analysis of the conjoint exercise revealed four broad groups of people when identifying opportunities for data science. Further exploration of the opportunities for data science should consider how best to engage with each of the following groups:

- **'data adopters'** (23% of adults) who support using data science for research purposes and see the value in how individual level data can generate better insight;
 - **'data adapters'** (28% of adults) who respond best to uses which improve services for individuals and use of non-sensitive data;
 - **'data pragmatists'** (27% of adults) who are more ambivalent in their views, wanting government to explore new ways of using data but are most comfortable using data for high-level statistics rather than advanced data science; and
 - **those who are 'data wary'** (22% of adults), who apply caution to the principle of data science, based on concerns around privacy and effectiveness or a desire for further information.
- **Attitudes towards government use of data science also differed based on individual experience of sharing data either with government or the commercial world.** Those who regularly interact with a number of different government services were often quicker to see the benefit of policy objectives (based on their experience of services); furthermore, those who are more used to sharing data through digital interactions were often quicker to see value in the concept of data science. However this is often accompanied by a greater ability to identify potential risk and therefore does not guarantee greater support for individual projects.

1.5 Key findings: towards a revised ethical framework

- **On the whole, participants agreed that the six principles of the draft ethical framework covered most of the key elements they thought should be priorities for people working in data science.** Participants' spontaneous priorities for the framework focussed mainly on transparency and standards. This included a broader request to be open and honest about government's approach to data science in general, as well as having accountability in place for specific projects. Participants felt that data science within government should be run with integrity and impartiality, with high standards in governance, oversight and regulation that are relevant and up to date. Though this would not in itself lead change the boundaries of what is considered acceptable, adherence to such standards would increase trust in how government uses data science.
- **Based on the findings from the workshop, three areas of the current principles in the draft framework would benefit from further clarification or elaboration.** These are:
 - (1) **Transparency** – it would be beneficial to have more guidance on how people working in data science should be transparent in practice. This includes how they communicate data science to the public as part of open policy making. Building on the current framework, good practice on gaining public trust and confidence would be helpful.
 - (2) **Outcomes** – the risk and proportionality of 'outcomes', both intended and unintended needs to be clearly recognised in more detail within the ethical framework, especially given the impact that this appears to have on the public assessment of whether data science opportunities are suitable.
 - (3) **Robust data science models and analytical approaches** – further clarity around robust data science models are defined, what they look like, what approaches are used, and what standards they are subject to would be beneficial. This is especially important when thinking about how best to be transparent and communicate data science to the public.

2. Introduction and methodology

2 Introduction and methodology

2.1 Background

Conservative estimates suggest 2.5 quintillion bytes of data are created every day. We generate records of data in almost everything we do from visiting health services and using travel cards, to shopping online and communicating through our smartphones. The number of digital interactions we make has increased rapidly, as has the number of devices we use that record and transmit data such as smart energy meters and parking sensors.

The proliferation of data generated every minute is striking, and suggests that the volume and rate at which we generate data is only likely to grow. To take Twitter as an example, in 2011 Twitter users sent over 100,000 tweets per minute, this rose to 277,000 per minute in 2013 and to 347,222 per minute in 2015. New services continue to leave digital footprints, such as the 694 rides taken by Uber customers or the 1,041,666 videos played on Vine every minute in 2015.⁵

Although much of the data we generate is unstructured and cannot be analysed using conventional tools, the potential for insight is vast. The ability to analyse these new forms of large and complex data through data science can not only help businesses understand current and potential customers or personalise its offer to individuals, but it also presents an opportunity for government to improve policy and public services. To name but a few examples, data science can help plan infrastructure, inform research, identify trends in population characteristics, spot future surges in access to services, and help keep citizens safe. Government therefore has a responsibility to explore the potential use of data science, and improve its understanding of how to maximise the benefits whilst mitigating against potential risks.

The Government Data Science Partnership was created to help meet this objective. It includes the Government Digital Service, Office for National Statistics and the Government Office for Science and acts to promote the use of data science across Government. There are four main strands to the Government Data Science Partnership work on data science:

- Alpha projects: to prove concept and identify where data science can help
- Skills & capability: to work out what skills we need in future, how best to get them and how to integrate with existing professions

⁵DOMO, Data never sleeps,

<https://web-assets.domo.com/blog/wp-content/uploads/2013/07/DataNeverSleeps.jpeg>

<https://www.domo.com/blog/2014/04/data-never-sleeps-2-0/>

<https://www.domo.com/blog/2015/08/data-never-sleeps-3-0/>

- IT & other barriers: to identify and develop solutions to IT and data gaps and barriers
- Ethics & transparency: to proceed through a transparent and ethical approach

Under the fourth strand of the partnership, the Government published a draft framework for Data Science Ethics in December 2015.⁶ Ipsos MORI was commissioned to undertake a public dialogue to inform the further development of the ethical framework.

2.2 Objectives

The aim of this research project was to test how the Government Data Science Partnership can maximise the opportunities of data analysis in an appropriate way. The objectives of the project were to:

- Explore, understand and report on the **opportunities for data science projects** within government (including what type of data science projects, the public benefit, the type of data used, privacy risks) the public think are appropriate and how these should be overseen;
- Use this insight to **inform an ethical framework** for departments to use through the detailed analysis, reporting and use of the insights generated by the dialogue;
- Develop and use a number of **case studies** in the dialogue process to enable participants to explore the ethics of specific data science projects;
- Explore, identify and report on participants' views on future **oversight and engagement**;
- Create and develop an **online survey** to create robust qualitative evidence on what the public thinks makes government data science projects appropriate.

The project was overseen by GDS, Sciencewise, GO-Science and ONS.

2.3 Methodology

The project consisted of three methods of public engagement:

- A series of **deliberative public workshops** to develop qualitative insight into public opinion on the appropriate use of data science. Reconvened workshops were conducted with 88 people across London, Taunton, Sheffield and Wolverhampton.
- Recruitment for and analysis of an **online survey** to develop robust quantitative insight into public opinion on the appropriate use of data science. An online survey of 2,003 people aged 16-75 in Great Britain was conducted between 24th February and 7th March 2016 using the Ipsos MORI Access Panel. The survey contained a conjoint module to identify the underlying factors driving attitudes and decisions about government use of data science.

⁶ Data Science Ethics, 8th December 2015, <https://data.blog.gov.uk/2015/12/08/data-science-ethics/>

- Use of the results of the online survey to develop a visual **engagement tool** to engage a wider audience in a public debate around data science. Based on the conjoint analysis, a draft online engagement tool has been developed in partnership with Codelegs. This is likely to be launched in May 2016.

Approach to deliberative workshops

Participants in the general public workshops were recruited face-to-face, on-street by specialist Ipsos MORI qualitative recruiters. A mix of urban, semi-urban and rural locations was chosen to ensure good geographical representation, and recruitment quotas were set to ensure that overall people of a range of ages and from a variety of ethnic and socio-economic backgrounds took part. Quotas were also set based on attitudes to government, digital literacy and number of interactions with public services.

A deliberative workshop approach was taken due to the complex nature of the issues around data use and reuse, and the low levels of awareness and understanding with which many people approach these issues. A workshop is an ideal, **open environment** that gives people time and space to learn new information, ask questions, change their minds and develop their views with others like them. Workshops also allowed sufficient time to explore a larger number of variables via case studies and other stimuli so that participants were able to see how data-sharing currently operates in the healthcare and research system.

Approach to the online survey

As part of the online survey, respondents were asked to complete a conjoint exercise which asked them to choose between multiple data science projects with different methodologies. 'Conjoint analysis' was conducted in order to identify the different principles that are important to people when faced with different data science projects.

This analysis involved asking respondents to imagine themselves as a part of a team in government responsible for solving problems using data science techniques. Respondents were then presented with different scenarios where data science could take place. There were five different scenarios, of which a participant would only see two. They covered the following areas:

- using data science for counter-terrorism;
- using data science for catching fare dodgers;
- using data science to provide personalised career advice;
- using data science to improve transport services; and
- using data science to provide advice on healthy foods.

Within these scenarios, respondents had the option of two different data science projects to choose in order to achieve the objective. Respondents were also given the option to take no action and told that this means the

aim would not be achieved if that response is chosen. An example of a type of screen that a respondent could see is below:

There are a number of different ways in which you could design the project. Here are two options which have different levels of effectiveness, but also different areas for concerns. Which of these two options, if any, would you choose to do to solve this problem?

	OPTION A	OPTION B
What type of data will be used?	Traffic and transport use	Criminal records
Is this information about individuals?	Not related to people or humans	Grouped into a large number of people, for example gender
How many people does the data include?	None, the data is not related to people	Everyone in the UK
How will the government use this data?	To research groups or areas which need more resources	To target specific groups of people (e.g. people in a certain area)
Who is making the the decision? Is it automatic?	Computer program makes a recommendation for action, but a human makes the final decision	Decisions are automated based on a computer program with no human control if something goes wrong
How clear is the decision? Does it tell you how it has been made?	The way the computer makes decisions is fully transparent and published for anyone to see	The way the computer makes decisions and learns is not known by staff working with it, only by the creators of the program
I would rather take no action, EVEN THOUGH THIS MEANS THE AIM WOULD NOT BE ACHIEVED		

Within each option were six different 'attributes' – categories of key elements of the project – such as the type of data that would be used, and whether the data includes information about individuals. Finally, within each 'attribute' were several 'levels' that represented a spectrum of the possible options within each attribute. Levels were randomly rotated to ensure that every combination of different scenarios was seen by a substantial proportion of the sample. For instance, within the attribute 'What type of data will be used?', a random selection of respondents would see 'Traffic and transport use', while a different selection would see 'Travel or store card data'. The subsequent analysis determines which attributes are the most important in the decision making process and, within that, which variables impact that attribute the most.

The following table (split across two slides) is designed to convey the different options that respondents might see. Green boxes represent the 'attributes', whereas purple boxes show the possible 'levels' that could be seen next to that attribute.

Attribute	Levels
What type of data will be used?	<ul style="list-style-type: none"> Sensitive personal information Travel card or store card data / Annual income Social media posts Criminal records / Tracking phone location Traffic and transport use / Live job vacancies / Gym memberships Official statistics
Is this information about individuals?	<ul style="list-style-type: none"> Individuals can be identified Individual data will be used, but people will not be identified Grouped into a large number of people Not related to people or humans
How many people does the data include?	<ul style="list-style-type: none"> Everyone in the UK Everyone in a specific area of the UK Only people who have used a service Only people who have agreed to have their used in this way None the data is not related to people
How will government use the data?	<ul style="list-style-type: none"> To research groups or areas which need more resources To target specific groups of people To target individuals to contact

Ipsos MORI Social Research Institute

Document Name Here | Month 2018 | Version 1 | Public | Internal Use Only | Confidential | Study Confidential (DS/ET CLASSIFICATION)

Attribute	Levels
Who is making the decision? Is it automatic?	<ul style="list-style-type: none"> Decisions are automated based on a computer program with no human control if something goes wrong Decision are based on a computer program with humans able to intervene if it goes wrong Computer program makes a recommendation for action, but a human makes the final decision A human uses computer programs to understand different options for action, but the human makes a recommendation and decision themselves
How clear is the decision? Does it tell you how it's been made?	<ul style="list-style-type: none"> The way the computer makes decisions and learns is impossible to determine The way the computer makes decisions and learns is not known by staff working with it, only by the creators of the program The way the computer makes decisions and learns is known to the staff working with it The way the computer makes decisions is fully transparent and published for anyone to see

Ipsos MORI Social Research Institute

Document Name Here | Month 2018 | Version 1 | Public | Internal Use Only | Confidential | Study Confidential (DS/ET CLASSIFICATION)

This analysis is intended to provide an understanding of whether, for example, **What type of data will be used?** exerts more influence on decisions than the question **How many people does the data include?** It is possible to look deeper into this attribute and see whether **Traffic and transport use** data is considered more acceptable in a specific data science scenario than **Travel or store card data**. Further technical details of the conjoint analysis can be found in the annex of this report.

Developing materials

Members of the Government Data Science Partnership convened an Advisory Group to provide a sounding board and critique for the design of the study and conjoint modules. Other stakeholders outside this group also contributed to the development of materials. We have listed membership of the Advisory Group and other contributing stakeholders in the annex.

The Advisory Group met at the start and during the study to help scope design and to hear and comment on initial presentation of findings to refine this report. The group also contributed to the development of materials by reviewing the case studies which formed a key part of both the qualitative and quantitative study.

The GDS Partnership and Ipsos MORI collaborated to develop a set of case studies. Five detailed, in depth case studies were designed to introduce the concept of data science and its potential uses by government during the first workshop events. A further nine hypothetical abstracts were developed for the second round of workshop events which reflected hypothetical uses of data science by government, thereby allowing us to explore the full range of issues with the public across a spectrum of most ethically acceptable to least ethically acceptable.

Five scenarios of potential policy areas for data science were also presented to respondents of the conjoint exercise. The annex to this report contains further details of all stimulus used as part of the deliberative process.

The tables below shows a summary of the case studies used at the workshops.

In-depth case studies (shown at event 1)

A: WHERE IS SAFE TO EAT? (in collaboration with the Food Standards Agency)

This case study explored the possibility of using reviews on restaurant review sites as an indication of hygiene ratings that could be used in inform FSA inspections.

This project tests the following data science elements:

- Web scraping
- Predictive tools

This project tests the following ethical considerations:

- Robustness of data sources (biased or non-representative data)
- Using data beyond original purpose
- Level of automation
- Transparent algorithms

B: WHO HAS CONFIDENCE IN THE COURTS? (in collaboration with the Ministry of Justice)

This case study explored the possibility of using social media data to better understand experience of the UK court system.

This project tests the following data science elements:

- Social media analysis for overall trends

This project tests the following ethical considerations:

- Using social media to identify people
- Robustness of data sources

C: WHERE ARE THE POPULATION? (in collaboration with the Office for National Statistics)

This case study explored the possibility of using mobile phone data to inform population estimates.

This project tests the following data science elements:

- *Use of anonymised mobile phone data*

This project tests the following ethical considerations:

- *Using sensor data/IoT and real-time data*
- *Linking data sets*
- *Inferring information*

D: WHO LIVES AT YOUR HOUSE? (in collaboration with the Department for Work and Pensions)

This case study explored the possibility of using commercial and government held data to help tack benefit fraud.

This project tests the following data science elements:

- *Use of non-government held data to identify people*

This project tests the following ethical considerations:

- *Unfairly adversely affecting someone*
- *Transparent algorithms*
- *Robustness of data sources*
- *Consent*
- *Acquisition of data (e.g. payment, enforcement)*

E: WHAT'S BEST FOR ME?

This case study explored the possibility of use of services to tailor those offered by government e.g waste disposal services.

This project tests the following elements of data science:

- *Segmentation*
- *Tailoring*
- *Combining multiple datasets from gov. and non-gov*

This project tests the following ethical considerations:

- *Consent*
- *Level of automation*
- *Repeating historic bias*
- *Limiting choice*
- *Linking datasets*

Hypothetical general purpose and method case studies (shown at event 2)To decide what kind of pension benefits will be right for you

- *Linking up data from across government*
- *Exploring past government data to predict what is likely to be best for individuals*

To decide what type of employment support is best for you

- *Looking at commercial datasets for past employment and training data*
- *Using past government service data to predict what is likely to be best for individuals*

To identify people leading unhealthy lifestyles

- *Looking at commercial datasets to understand behaviours*
- *Analysing individual health records*
- *Analysing aggregated education and health records*

To identify people living in UK illegally

- *Using individual-level sensor data*
- *Matching-up individual health records with visa data*
- *Looking at individual data on 'overseas students'*

To identify deliberate fare dodgers

- *Looking at mobile phone GPS data*
- *Monitoring smart cards usage*
- *Comparing non-sensitive data on footfalls with station ticket sales*

To identify terrorists

- *Using algorithms to explore content of emails*
- *Looking at IP addresses to observe behaviours*
- *Looking at online postings on public forums e.g. social media sites*
- *Responding to 'tip-offs' of suspicious behaviour*

To identify motorists who break the speed limit

- *Using individual-level sensor data to monitor speed*
- *Using GPS mobile phone data to identify average journey times*
- *Exploring historic data to identify groups or individuals likely to speed*

To understand sexual preferences

- *Looking at identifiable Twitter data to find groups or individuals who are gay/lesbian/bisexual*
- *Looking at aggregated data collected using equality forms by commercial companies*
- *Exploring historical census data to understand individuals' preferences*

To understand future school age population

- *Looking at social media site data to understand where families with kids nearing school age are moving to*
- *Linking GPS mobile phone data to contractual data to understand population movement*
- *Linking aggregated education records with health (vaccination) records*

To predict future levels of cancer across UK

- *Looking at commercial data for online sales of books with health focus*
- *Exploring personal/aggregated supermarket data*
- *Analysing health records*

2.4 Interpretation of data

Results from the online survey are based on all respondents unless otherwise stated. Please note that where percentages do not sum to 100, this may be due to respondents being able to give multiple responses to a question or computer rounding. An asterisk (*) indicates a percentage of less than 0.5% but greater than zero.

A dash (-) indicates a percentage of zero. The data has been weighted to be representative of gender, age, region and working status.

When assessing insight generated by the qualitative workshops, the following note may be helpful. Qualitative research approaches (including deliberative workshops) are used to shed light on *why* people hold particular views, rather than how many people hold those views. It is used to explore the nuances and diversity of views, the factors which shape or underlie them and the ideas and situations in which views can change. The results are intended to be illustrative rather than statistically reliable. Given the qualitative nature of the data collected from the workshops, this report aims to provide detailed and exploratory findings that give insight into the perceptions, thoughts and feelings of people, rather than statistical evidence from a representative sample.

It is not always possible in qualitative research to provide a precise or useful indication of the prevalence of a certain view, due to the relatively small number of participants generally involved (as compared with the larger respondent bases involved with quantitative studies). So, the views of proportions of the qualitative group should not be extrapolated to the population at large. Sometimes, ideas can be mentioned a number of times in a discussion, and yet hide the true drivers of thoughts or behaviours; or a minority view can, in analysis, turn out to express an important emergent view or trend. The value of qualitative work is to identify the issues which bear **future investigation**.

Therefore we use different analysis techniques to identify how important an idea is. The qualitative report states the **strength of feeling** about a particular point rather than the number of people who have expressed that thought. Having said this, is it sometimes useful to note which ideas were discussed most by participants, so we also favour phrases such as "a few" or "some" to reflect views which were mentioned infrequently and "many" or "most" when views are more frequently expressed. Where views apply only to a subset of participants, e.g. participants in Taunton, we have highlighted this in the text, as this may indicate differences by rurality, for example. Any proportions used in our qualitative reporting (e.g. a "couple of" or "a few" participants), should always be considered indicative, rather than exact.

Verbatim comments have been included in this report to illustrate and highlight key points, i.e. those views either shared by a large number of participants or reflecting the strong views of a smaller subset. Where verbatim quotes are used, they have been anonymised and attributed by location and group/workshop-type (e.g. General Public, Dundee).

The qualitative findings complement the quantitative findings. While quantitative and qualitative methodologies are inherently very different, they were designed to complement each other in answering the same research objectives. Workshop participants went on a much more substantive journey through the day and their views were nuanced. In the qualitative work, there was more scope for sharing information about data science context and the details of specific projects. In the quantitative survey there was less opportunity to provide respondents with background information or indeed for them to truly deliberate; however we were able to test underlying attitudes (rather than stated attitudes) through the conjoint exercise, and run split-sample tests to establish the impact of different framing on attitudes.

2.5 This report

This report explores the opportunities for government use of data science across six further chapters:

- **Chapter 3** explores the context for discussing data science, including awareness and underlying initial assumptions of how data is generated, how it is used in data science, and how government is currently analysing data. This will help policy makers and data scientists assess the current perceptions held by members of the public.
- **Chapter 4** reflects on how members of the public engage with the concept of data science, and discusses the opportunities for future engagement strategies.
- **Chapter 5** explores the process and principles applied by members of the public when assessing opportunities for government use of data science. This will help policy makers and data scientists evaluate whether the benefits of specific projects are likely to be perceived to outweigh the risks.
- **Chapter 6** differentiates between key groups of public attitudes. It identifies key clusters of opinion and explores the different public mindsets that should be considered by policy makers and data scientists when scoping or designing projects.
- **Chapter 7** identifies the key overarching considerations that should shape an ethical framework for data science, informing the further development of the current draft published in December 2015.

Alongside this report, further information about the recruitment, stimulus and questions asked for both quantitative and qualitative components can be found in a separately published Annex. The Annex also includes a topline of responses to the online survey. Other outputs from the project include a video with interviews of participants at the workshops, and an online engagement tool 'Data Dilemmas' which asks users to design the components of their own data science project.

2.6 Acknowledgements

The project team at Ipsos MORI would like to thank all who provided help and support in the design and delivery of this research, in particular: the core project team at the Government Digital Service (Cat Drew and Madeleine Greenhalgh); staff from Sciencewise who assisted in the running of the project; all members of the Advisory group, and the individual stakeholders who attended the events as experts, and assisting in the smooth running of the day.

3. Data science in context

3 Data science in context

This chapter introduces the beliefs, views, and assumptions people hold around data and government use of data. It presents the context for exploring opportunities for data science discussed in later chapters.

Summary

Despite being comfortable with the concept of sharing personal information to access services, participants were unaware of the volume and rate at which they generate data in numerous sources across their everyday lives. Furthermore, there was also little knowledge about how individual data records can be anonymised or aggregated to create larger datasets in order to conduct analysis and generate insights. The term 'data science' itself was very unfamiliar.

Consequently, participants were often unsure and cautious about how data is, or could be, collected and used by government for policy-making and planning services. With very low awareness of how data science works in *any* context, many initially struggled to see the value of using computer analytics in this field. This led to some scepticism and meant that many adopted a tentative position towards government adopting data science methods.

However, across both the workshops and online survey, participants were broadly willing to support the need for government to find new ways of using data in general, recognising the potential uses and relevance that the huge volumes of data we now generate and collect have. Concerns arose when considering the power and significance of the decisions that government could potentially make using data science. Many spontaneous responses to data science focused on the issue of data security. While not exclusively an issue about data *science*, this highlights the fact that the public do not separate concerns about data from data science.

Initial reactions to the subject of data science reveal that the way people digest and react to the topic is shaped by several factors, namely their relationship to and trust in government, their use and attitude towards data and technology and their awareness of the data collection that currently takes place by government and others.

3.1 Generating data

Early stages of the public workshops explored participants' initial level of knowledge of data and data science. Overall, participants across the workshop events had moderate awareness of the data they were sharing. In general, they were good at identifying instances when they actively gave out personal data, but were often surprised to learn of circumstances where they were unconsciously generating data. Participants' awareness of data generation could be broadly be categorized into three levels:

1. **Personal interaction** – *Actively and consciously* giving information. Participants mentioned many instances when they had given out personal data knowingly – such as their name, date of birth, email address, postal address and credit card information – in order to access services. Common examples included signing up to websites, creating bank/other accounts or making an online purchase, and accessing health services.
2. **Everyday low-interaction sharing** – Aware that data is *probably being generated*, but *not consciously thinking* about it in terms of 'data'. In these cases, participants were aware that they were generating a record of some sort, but often as a by-product of another conscious action and/or mostly perceived as pseudo-anonymous and therefore of less significance. For example, using a travel card, or making a phone call on their mobile phone. In each of these examples, participants did not perceive the information to be personal and were largely unaware of the potential for data to be anonymised and/or aggregated to create larger datasets and generate broader insights. Social media posts were generally perceived in this category, as participants did not associate sharing photos or updating their status as generating 'data'. Web browsing history was a more familiar example, as most participants recognised that online adverts could be based on their recent browsing activity.
3. **Unknowingly sharing** – *Complete unawareness* of data generation. Participants were less familiar with data that is generated without conscious intent, or for purposes outside its primary or a closely related use. For example, sensors able to track whether car spaces were free, energy smart meters, or mobile phone location data that could be used to analyse traffic movements. In these instances, many participants were not consciously aware that various aspects of their behaviour were being recorded as data in order to generate insight either in real time or at a later date.

As a result of the deliberative process, many participants reported that their awareness of the data they generated had greatly improved:

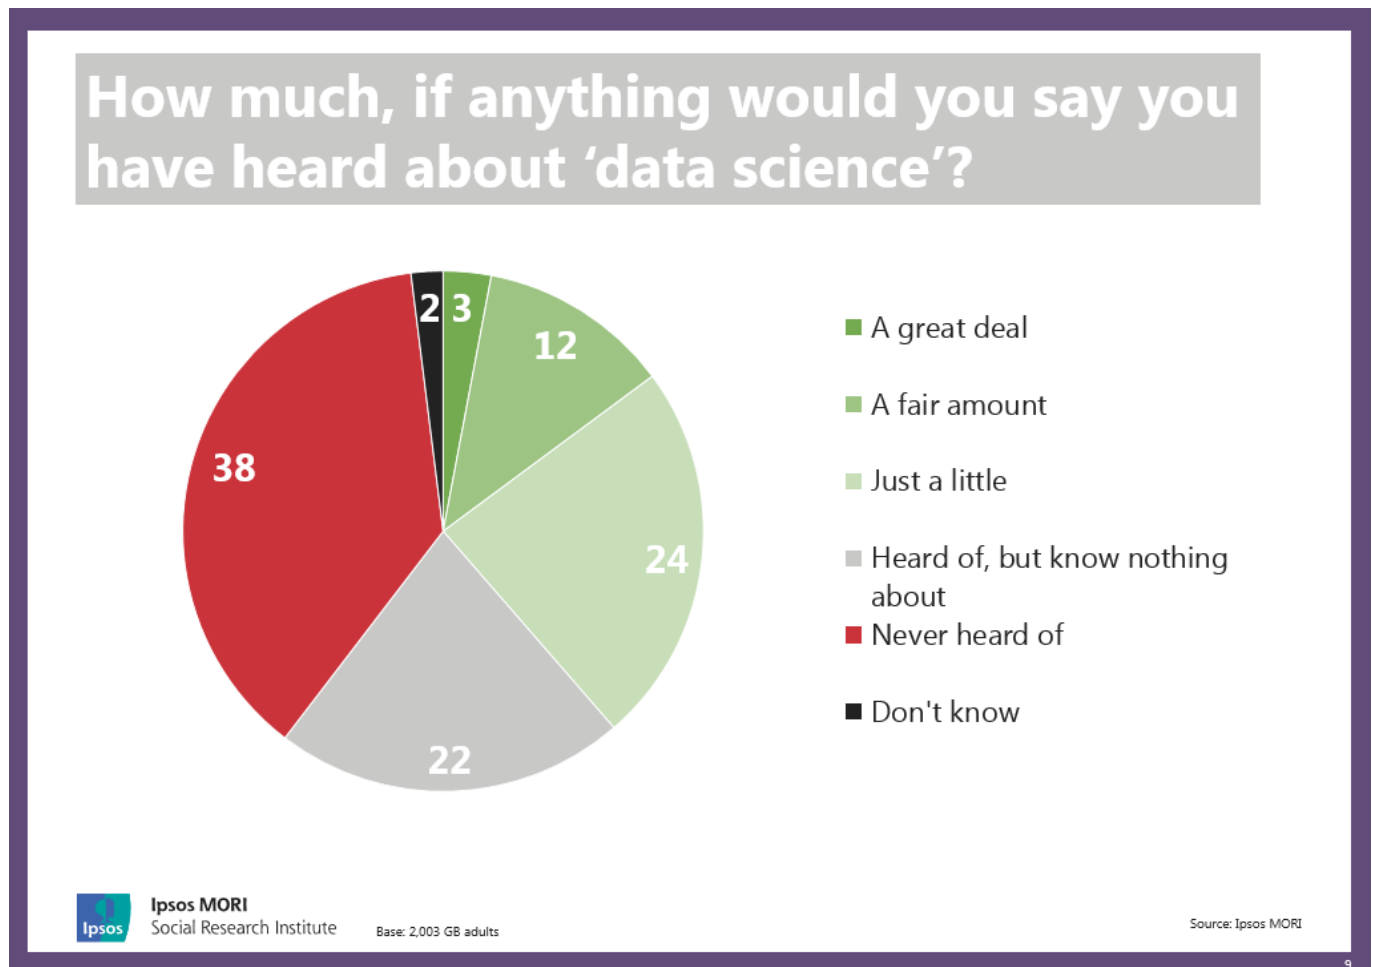
Before these sessions I had no idea how much data I was giving, how much I was sharing and how much what I did online affected what happened online.

Taunton, event 1

After the first event, workshop participants had a homework task where they were asked to write down: a) when they notice themselves creating/sharing/giving data and b) when they think data is being collected about them and why. Between the two events, participants were much more able to give examples of the second category, '*everyday low interaction sharing*', but they continued to struggle to give examples of how data could be generated without intent.

3.2 From data to data science

Awareness of the term 'data science' was low. Among those completing the online survey, only 15% had heard a great deal/a fair amount of data science, a further 24% have heard just a little. Broadly in line with workshop participants, more than one third had never heard of the term before.

Figure 3.1: Awareness of data science among general public

Beyond awareness of the term data science itself, participants held a series of initial assumptions about data analysis more broadly. These were often borne out of a lack of awareness about how the mechanics of data science worked:

- 1) Firstly, a number of participants in the workshops found it difficult to understand that the same data interaction could generate many different types of datasets. For example, that personal information generated through visiting A&E could be used to generate one of a large number of aggregated records to measure waiting times, or anonymised to inform research about the characteristics of a health condition. Participants often assumed that individuals could be identified in all types of datasets, and thus privacy was of paramount concern in every case.
- 2) Some participants also found it difficult to believe or trust that computers could make better decisions than humans. These participants were quick to point to challenges of data accuracy and often overestimated people's capability of making rational decisions. They had a preference for spending more money on human-based decision making rather than have an algorithm make decisions.

My judgment is more accurate than a computer using past data.

Sheffield, event 2

3) Many participants were ambiguous about the level of control and automation that can or should be given to a computer during a data science project. This fed the expectations and demands that they made of the humans involved in data science and led people to consider both the merits and drawbacks of human oversight and involvement. For instance, some workshop participants were fearful that computers might remove the need for human judgement entirely and they assumed that power for decision making would automatically be transferred to a computer algorithm. Contrastingly, others recognised the reality that the need for human design and interpretation during the data science process also brought downsides as it meant that there would always be an element of subjectivity; the use of computers does not necessarily remove bias or guarantee any more objectivity than if humans alone were involved.

Human beings have to instruct the computer. If we take a step back, who creates machines? Humans do. There will always be an element of bias. If I'm passionate about something, no matter how objective I want to be. If we are looking for certain trends, we are inputting data. It's all good getting correlations and trends, but we will interpret as we see fit. We [need to] have measures that we put in place to check things are being done properly.

Taunton, event 1

4) Many participants were also sceptical of the link between correlation and causation, and were most cautious towards patterns in datasets that appear unrelated. For example, in the examples below participants called into doubt how the amount and type of programmes that someone watches on Netflix could be a suitable predictor of health; or how the sales of books related to cancer on Amazon could be a suitable predictor of an increase in cancer cases. Correlations in data science can offer new insight, but do not always demonstrate causality. Where correlations existed between big datasets in the scenarios presented, participants were often quick to point out the exceptions to the rule which, as they saw it, rendered the whole approach 'un-workable'; moreover, they found it difficult to grasp how in some circumstances, non-causal correlations could still be helpful in data science.

Just because you're watching Netflix, it doesn't mean you are unhealthy

Sheffield, event 2

I could be searching for family members, or research. But it could be an indicator, but that figure in itself is no more than a spike. Doesn't mean that more people have cancer. Just that there is more interest.

Taunton, event 2

5) Finally, participants were cautious of how individuals could be clustered together into groups or segments using data science techniques. This was often based on an unease about being placed in a segment that they did not associated themselves with, or again, when associations were made between characteristics that appeared unrelated.

It's taking away from humanity, we're all just a number. It's generalising, putting people into labelled groups. But it doesn't work like that with humans. We're unique.

Sheffield, event 1

Despite these initial assumptions, most participants responded positively to the broad motives and the potential of data science as demonstrated across the two workshops. This did not guarantee approval of specific projects, all of which came under political, ethical and technical scrutiny; however many had understood that if applied correctly, data science could bring real value to society.

To be honest I have never heard about data science before, but having now come to these sessions, even though it's quite a complicated subject, what I have picked up is that it's trying to make Britain into a better place.

Sheffield, event 2

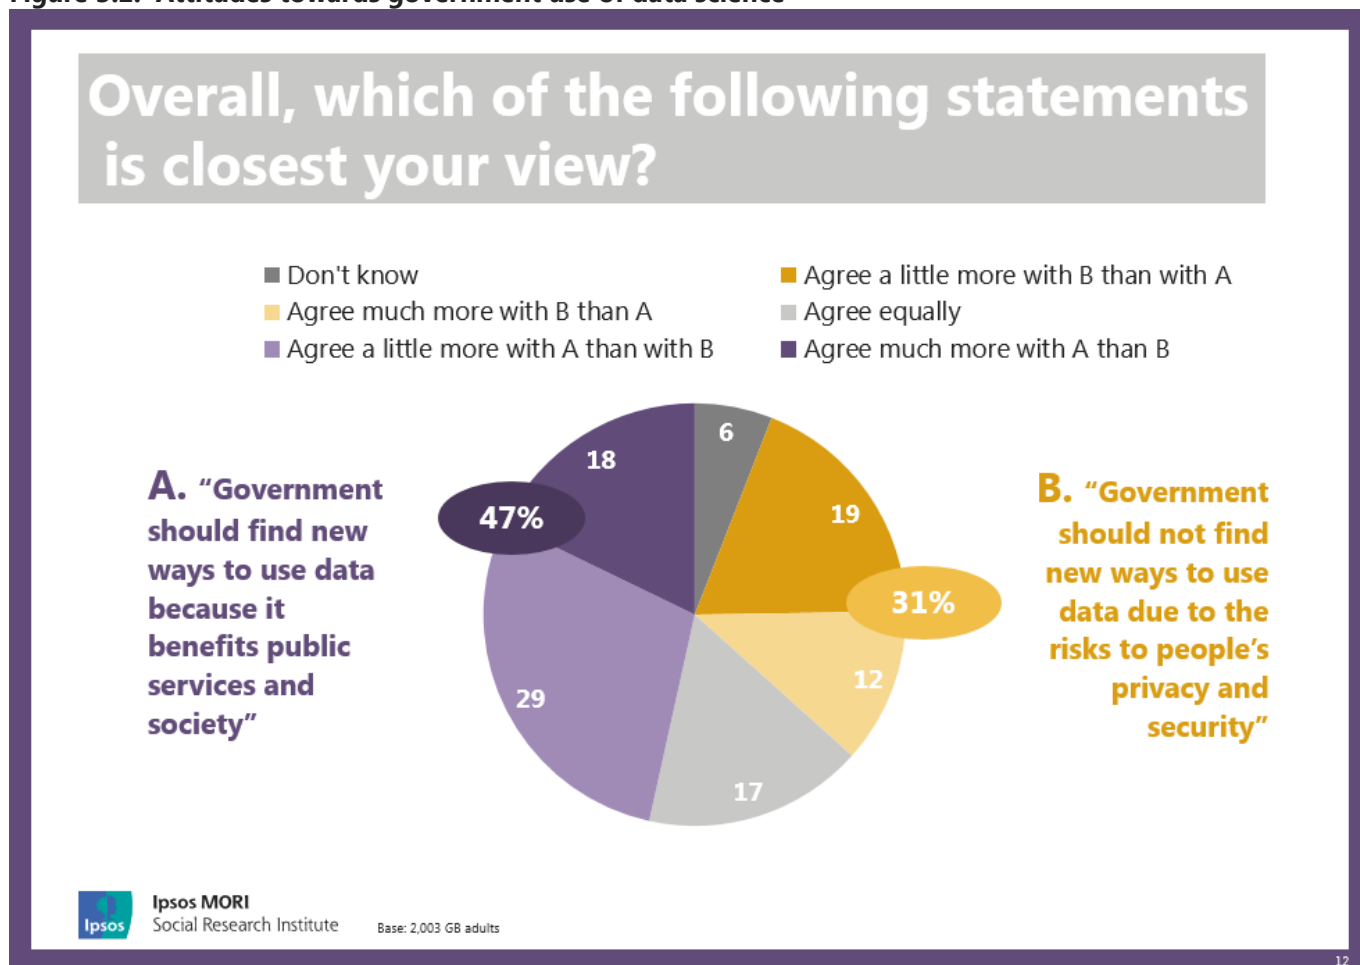
Chapters 5-7 explain further what the underlying reasons are for some of the public's caution and resistance to government's use of data science, which can be shaped by additional factors such as people's relationship to and trust in government, behaviours and attitudes relating to data and technology and awareness of current data collection by government and more widely.

3.3 Perceptions of government use of data science

The quantitative survey showed that opinion is mixed on the extent to which data science should be used by government. Close to half (47%) of adults are comfortable with government exploring new applications for data science; however just less than a third (31%) think that government should not explore ways of using data science due to the privacy risk.⁷ This sentiment was reflected in the workshop events, where most were immediately open to new ways of using data to benefit public services and society.

⁷ A quantitative experiment, described in section 4.2, demonstrated that further exposure to content made respondents more open to government using data science.

Figure 3.2: Attitudes towards government use of data science



Again, attitudes towards potential use of data science by government were based on a number of underlying beliefs and assumptions. During the public workshops, participants generally had a good understanding that their data is collected by private companies; such as when shopping online, when creating a Netflix account, when signing up to job-hunting sites, and when they share content on social media. It became evident that awareness of private companies' use of data is linked to an everyday interaction of sharing data. This was seen as having a clear financial purpose for private companies.

Yet awareness of how and why government organisations collect and use data was low. Most of the examples given by participants related to formal data collection, such as collecting data for the census or council tax. Most participants assumed that data is already widely shared between government departments. However, this assumption was accompanied by a mixture of positive optimism, or scepticism and concern about the volume and type of data known overall.

I'd imagine every government department is collecting data all the time otherwise how do they make decisions about how to do things better.

London, pilot

They [government] know everything about you. They're just coming up with new laws and stuff that aren't really protecting you. Just being watched, you don't feel free anymore.

London, 'high-tech' event 1

Where there was concern about government use of data, this was generated from a number of different attitudes and assumptions. These are important context for understanding why some participants were less open to opportunities for data science.

- 1) Some participants had very strong views that government was able to access information considerably beyond what is currently possible, for example how you voted in a general election. This assumption is in part due to the pictures portrayed on TV shows, where data is routinely seen as accessible to those investigating crime and other government departments. A few gave responses which gave the impression that what is seen on TV is actually what happens in real life.

They know everything! [from having watched CSI: CYBER]

Taunton, event 1

You get cop shows on TV and they say 'get their phone records' and you're not sure if this is possible or just for TV. It's creepy ... they can say 'he phoned X person at 10.45'.

Taunton, event 1

- 2) Most felt that data collection by private companies was out of their hands and something that they just 'have to accept'. Some felt that there was a trade-off and a clear benefit to sharing data, others felt uneasy that they felt they had little choice: for example, be on Facebook and accept that data is collected, or don't be on Facebook and be and lose contact with your friends; sign up to a service online and share some of your personal details, or do not get access to the service. This same sense of lack of empowerment was also present in views of the data relationship between citizens and the state, often in the context of privacy vs. public benefit.

Advantages and disadvantages, big brother is watching you. That aspect makes me worried. But better services I'm happy with those. If it helps health and education. Would not be happy with marketing if you're bombarded with phone calls and letters.

Sheffield, event 1

- 3) However, there were clear points of distinction between the commercial and government use of data and data science. The social role and legal reach of government generated quite different levels of concern about the potential for misuse of data in government rather than commercial hands. Participants cited that the government has the power to send someone to prison, stop their benefits, or close down their business. These powers presented not only a great opportunity for positive outcomes of data science, but also considerable risk where errors or misuse could occur.

In contrast, the risks of data science projects among private companies were perceived to be of much less importance: participants felt that in the "worst" case scenario, companies might try to make them spend more money, which they might not have thought about through targeted advertising. For some

participants, the concern was less about what would happen with the data collected initially or the project commissioned today, but the precedent it could set for future uses and the intentions of government in the future.

The problem is that you open up the floodgates and we haven't set the limits of what the government can use that data for. Who knows where that government can go. And it might change with everything government; they can do a lot of things. That's the problem.

London, pilot

Yes the government, they can be corrupt, so you never know what your data is being used for.

London, 'high-tech' event 1

- 4) A further concern among participants relates to protection of privacy and security. For those whom safeguarding data was of paramount concern, many cited a lack of confidence that government would be able to prevent data from being lost, hacked or otherwise misused. This scepticism was drawn largely from private sector examples, and concern that despite companies saying that they have security measures in place, there are ways to go around it - this could therefore happen if government stores data too. Others felt a sense of unease and knowing where their data will end up within government once it has been given to one department or service.

We like to think we are governed by the law, there is data protection, but is there? In terms of us giving consent, if government has the data, have we consented, does that hinder our freedom?

Taunton, event 1

I think my privacy settings are private but there's probably someone in China reading it and chortling.

Sheffield, event 1

3.4 Towards opportunities for data science

Yet, as participants developed an understanding of how data science works, they started to conceptualise opportunities for how government can use data science to help the public. Some participants suggested that there was a duty to explore the potential for data science within government; most saw the biggest area of potential in the planning of services.

We'd be the first to complain if the data isn't used. If there's info on population growth and there's not enough school places, the public would be the first to complain. Can be intrusive but sometimes a need.

Sheffield, event 1

They can plan services for trains and what they need to be providing for. There's no point just going blind and just thinking that something might be needed when in fact it isn't needed. So it cuts down, it'll cut down costs maybe...

Wolverhampton, event 1

Further discussion of the potential opportunities for government use of data science is contained in chapter 5.

4. Engaging the public in data science

4 Engaging the public in data science

This chapter explores the way in which members of the public engaged with the topic of data science, and draws lessons for how this could help design future engagement strategies.

Summary

Demonstrating the potential impact of data science through real life case studies was crucial to engaging the public on future opportunities for data science. Yet, although the introduction of case studies was key to moving beyond general attitudes towards data and the role of government in society, the presentation of further information did not necessarily affect the way in which participants assessed the risks and benefits of individual opportunities for data science.

Publication of the guidance offered by government on data science is more likely to reassure than concern people. However, although there is broad support for a wider discussion to help refine the opportunities for data science, the current lack of awareness and underlying assumptions means there is the possibility that some members of the public will become more concerned on learning of the multiple risks involved in any data science project.

4.1 Importance of case studies when discussing data science

Participants were shown a large amount of new and complex information about data science. As with any dialogue, the way this information is introduced needs to be managed carefully. Piloting workshop materials prior to the main dialogue workshops allowed us to improve the stimulus for future events and to draw lessons on how engagement on data science should be done in the future. It revealed the preconceptions that people have surrounding data and data science, and the gaps in understanding that if not filled can restrict detailed debate of opportunities, benefits and risks of data science design.

As outlined in chapters 5 and 7, decisions about the appropriateness of data science projects were not made in isolation from other attitudes towards the role of government and the characteristics of a good society. Objections to the application of data science within government were therefore a result of a number of different perspectives, including media literacy, trust in government, and in the policy that the project would contribute towards. This was made very apparent during the pilot workshop, where a few participants had such strong feelings of distrust towards today's political figures that they struggled to move beyond and discuss the potential benefits of data science.

For me it is very hard to say because with this particular government I have a trust issue.

London, pilot

Similarly, a questioning of the effectiveness of data science as a methodology emerged at the pilot. This was most evident in discussion of correlations, where participants emphasised that correlation is not the same as causation therefore questioned the validity of the data science project. Clearly, many members of the public are not as well-versed in data science as those working in the field and without accessible explanations of how data science works (e.g. why, in some circumstance a correlation between two sets of non-causal data could still be a useful tool for understanding) they are liable to pull data science projects apart and raise doubts about the value of using data science over other approaches.

It was these broader concerns that came to light in the pilot, often stalling more in-depth discussion of the issues and making it difficult for participants to engage with the detail of hypothetical future government case studies. The experience of the pilot therefore highlighted the importance of clearly communicating how data science has been used in the past, and what kind of benefits it has facilitated. Revised materials were developed that allowed people to discuss data science in a way that would make sense to them, including real-world examples from both the UK and further afield to show that data science has potential (see appendix for examples of revised stimulus). This freed the discussion up and moved people away from concerns about UK politics as well as providing evidence to show the fact that data science can lead to positive outcomes in practice. This empowered participants to feel able to comment on further examples of the potential opportunities for government use of data science.

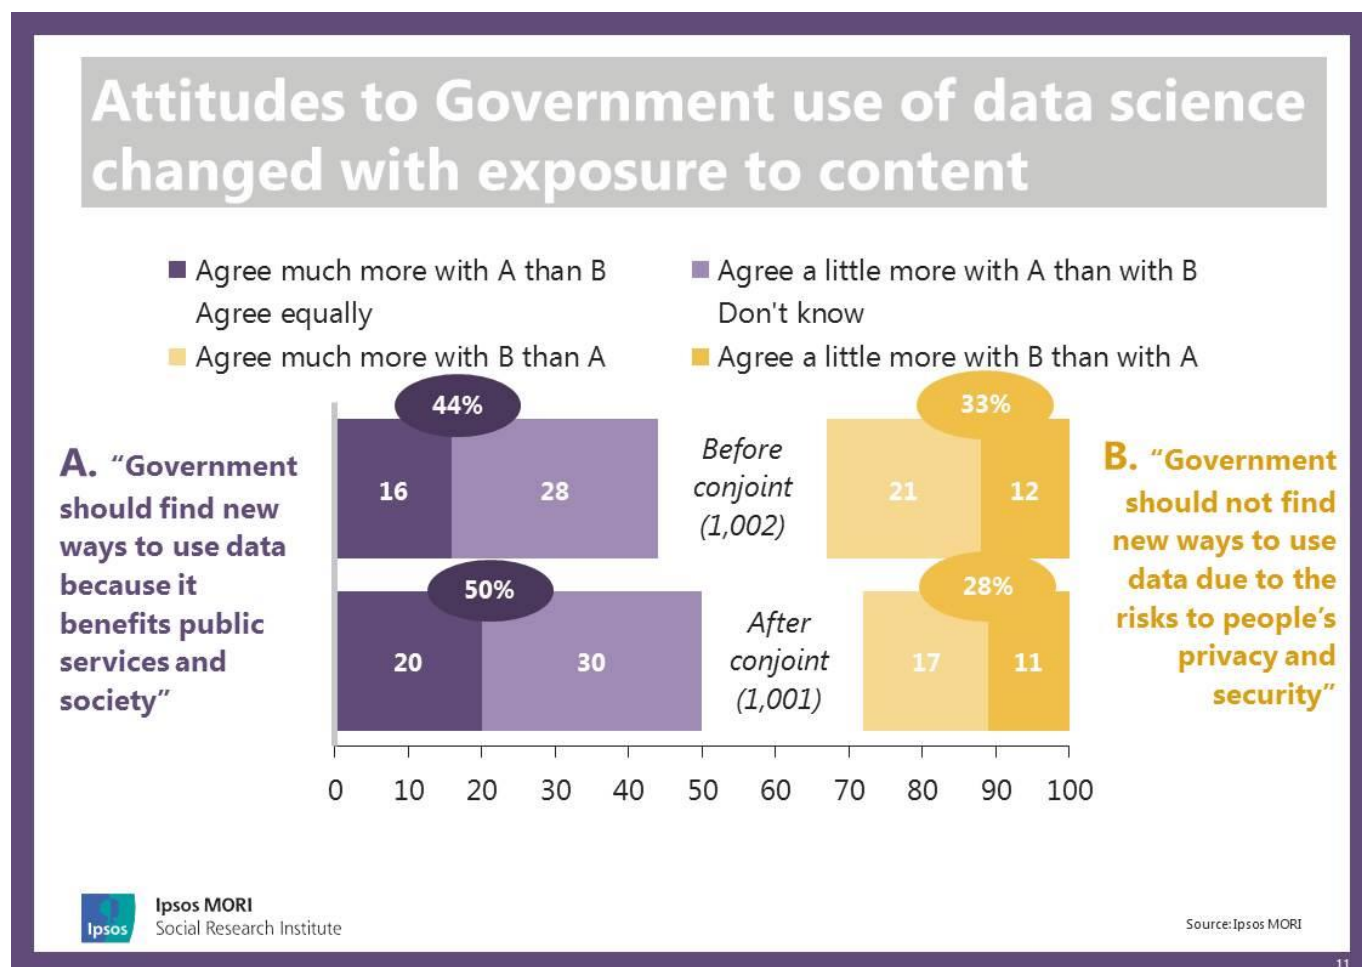
As the pilot showed, the risk of not providing enough concrete information on where data science has worked previously is that people fall back on worries about data in general or on their attitudes towards government.

It is also worth noting that some participants felt it would have been helpful to provide further examples of where data science has worked less well. It will therefore be important to reflect the challenges and potential pitfalls of data science whenever developing examples for future engagement with the public.

4.2 Quantitative experiments

Separately to the public dialogue elements of the project, we designed two experiments in the quantitative survey that were intended to explore this effect of giving participants more or less information about data science. These experiments were designed to understand the impact of information on attitudes towards data.

The first quantitative experiment involved asking a question of half the sample before the conjoint exercise, and half following the conjoint exercise (see Introduction for further explanation). The purpose of this was to understand to what extent showing respondents concrete examples of data science would impact their overall views of government using data. As the chart below shows, there was a significant change depending on whether this question was asked before or after the examples in the conjoint exercise.

Figure 4.1: Attitudes towards data science asked before and after conjoint exercise

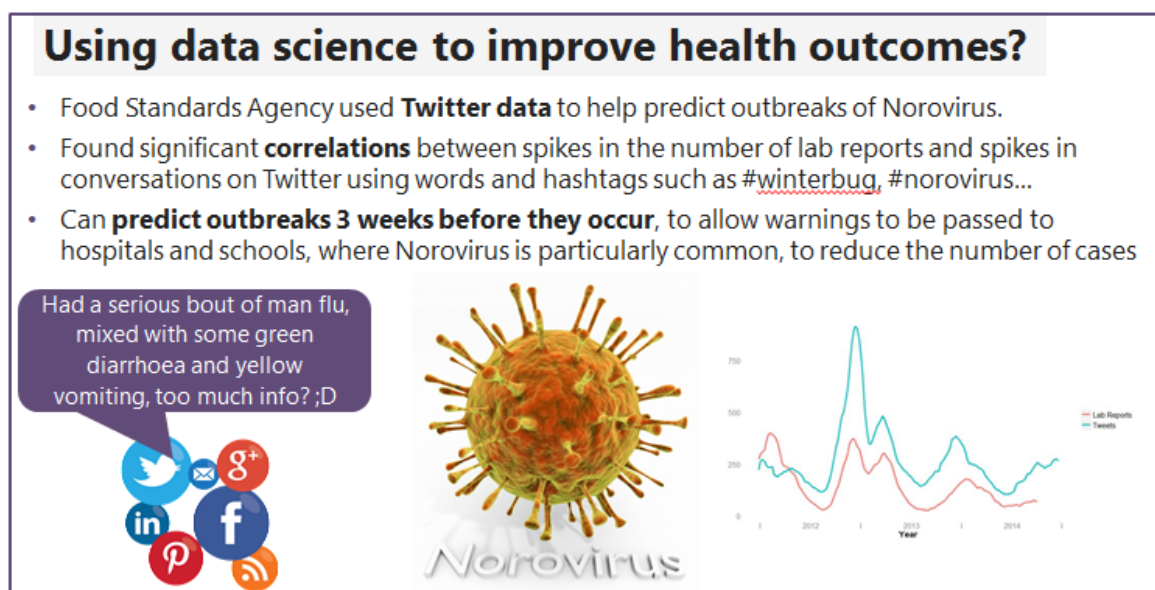
While under half (44%) agreed more that government should find new ways to use data when asked before the conjoint exercise, this proportion rises to half (50%) of those asked after the conjoint exercise. Conversely, the proportion who do not want government to find new ways of using data falls from a third (33%) before the conjoint to under three in ten (28%) after the conjoint.

This finding is consistent with our hypothesis following the London pilot. Attitudes towards data usage by government are affected by exposure to the benefits and opportunities provided by data science projects. It is therefore important for future public engagement that the tangible benefits resulting from data science are well-communicated, as it will improve their buy-in to data usage by government overall.

The story is not as simple as this, however. We also hypothesised that showing some additional information to some participants before they undertook the conjoint exercise would lead to a similar effect as the first experiment. As such, we showed 250 respondents an additional introduction screen before they undertook the exercise. We used a similar rationale to our changes following the London pilot, and so included a case study that showed a tangible example of what data science could look like, even using the same material for this case study as we used in the workshops, because it worked well as an explanation of how data science works:

Figure 4.2: Additional information shown to 250 respondents prior to them completing the conjoint exercise

Data science has been used by the government in different ways. For example, in the project below, the Food Standards Agency in the UK used social media data to predict outbreaks of the norovirus before lab reports can. In another project, IBM used anonymised mobile phone data in the Ivory Coast to suggest two new bus routes and an extension to an existing one – reducing travel times across the capital by 10%.



The result of showing this to respondents, however, **makes no discernible impact** to their choices in the conjoint exercise, where respondents were asked to judge the components and merits of individual projects. This means there is some nuance to the idea that more information makes the public more accepting of data use. Pulling in the understanding of the pilot, we postulate that this is due to the two experiments exploring different parts of the public's decision process about data science by government.

Providing information and examples to show how data and data science projects can lead to positive outcomes for the public, and allows people to “buy-in” to the initial concept that data can have a public benefit (as evidenced in our first experiment). Furthermore, as demonstrated in the pilot workshop, this buy-in is necessary to move on to a more in-depth discussion of the merits and drawbacks of distinct data science projects.

Unless the public benefit of data and data science is established early, public discussions will just focus around debating the validity of whether data has a public benefit and/or the intentions of government. It is, after all, a coherent view to agree about the public benefit of data usage and disagree with the terms of a data science project.

If businesses are using it...then the government should as well. If it saves money and improves services, then it is good. Like car registration numbers, I thought it was spooky but now I think it could be used for good. Using publicly available data effectively is good.

Sheffield, event 1.

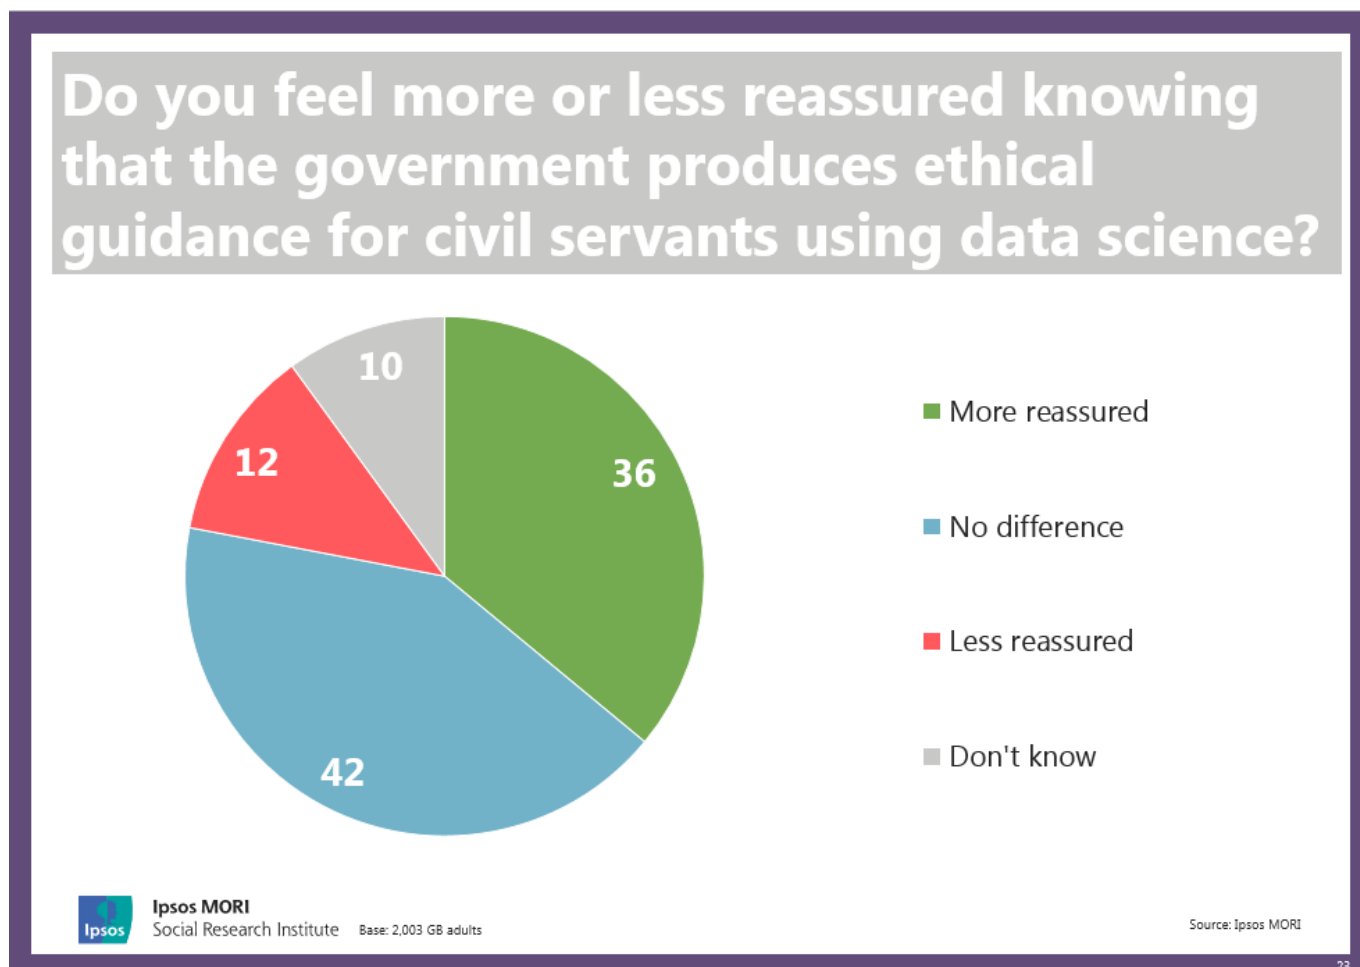
Once this potential public benefit has been accepted, there is still the outstanding question about whether providing more information affects how people weigh up the positives and negatives of specific data science projects and approaches. In the quantitative research, this is represented by our experiment of showing 250 respondents additional information to see if it changes their decision model. The experiment we conducted before the conjoint exercise provides some evidence that showing additional information does not seem to change the underlying risk calculations that the public make.

Providing extra information to participants can make them more likely to accept the public benefit of a policy area, but it does not necessarily change the next step in their thinking, which is to assess the risk of a data science project.

4.3 Future engagement

Participants at the workshops welcomed the opportunity to discuss the potential use of data science within government, and welcomed the work already undertaken as part of the draft Data Science Ethics Framework. In the quantitative research, when online survey respondents had completed the questionnaire (and conjoint module), more than a third (36%) stated that they feel more reassured to know that government produced ethical guidelines for civil servants using data science. Despite this, two in five (42%) said that it made no difference to them, a further 12% stated they felt less reassured.

Figure 4.3: Reassurance provided to the general public by ethical guidances of data science



The reassurance demonstrated through the online survey maps onto the demands for greater transparency and greater oversight stated by workshop participants (in chapter 7) when designing principles for government use of data science. Many participants at the workshop spoke of being *'more aware'* or having their eyes *'opened'* by the deliberative process and they were keen to have a more open and honest conversation between with public figures about how data could be shared and used in government-led data science projects.

I think from lot of discussion that we've previously had, I think not being educated about data science, which you're not going to be in your day to day life, is where your fears come from.
Sheffield, event 2

If we are going to have honest and open conversation, and we live in a society, it is a two-way process, it's a two-way contract. So if the Government want information from its citizens then we need to know how they are going to process that information, what purposes are they going to be using it for, so that we can be honest and open with them and work with them rather than against them.

Sheffield, event 2

Some participants, however, saw the practical burden and risks to effectiveness of requiring government to always consult with members of the public given the extent to which data science is applied in everyday life.

I don't think the Government should worry about it because we are opening up ourselves to data science all day every day... I think the government is being very polite!

Taunton, event 2

But others felt that duty held by government to its citizens justifies the need for additional engagement to ensure that its actions are just and can be held to account.

Their intentions are better than Google, but they need to look after the people of Great Britain. You need to feel that you trust them. They should be accountable for what they do.

Wolverhampton, event 2

Furthermore, as noted in the online survey, the experience of learning more about data science also has the potential to raise concerns among some people. Those who said they were less reassured by the production of ethical guidance were less likely to know about data science, less likely to already believe that government should explore new ways of using data, and felt less comfortable accessing services online. This suggests that multiple approaches will be required to address the variety of public sensitivities surrounding the issue.

Based on the experience of conducting the research, and on the feedback of participants, transparency should be a cornerstone of any future engagement with the public. As well as offering a transparent account of methods and benefits, future engagement on data science should:

- Account for low level awareness of data science, both of the data and techniques used, and of the detailed ethical trade-offs encountered in design;
- Use case studies of previous successful data science projects to demonstrate the potential value and impact of these methods;
- Account for the difficulty in separating attitudes to policy, data, government and data science technique;
- Give tangible examples of what any underlying principles for the framework mean in reality, for example using hypothetical scenarios that are familiar and easy for the public can relate to;
- Be aware that more education on data science opens up broad acceptability to government exploring opportunities for data science, but does not change the individual trade-off of risks and benefits for specific projects

5. Priorities and opportunities

5 Priorities and opportunities

This chapter will explore what **key considerations** the public apply when assessing government uses of data science, what their priorities are and where they see future opportunities.

In the dialogue workshops, participants spent extensive amounts of time discussing a range of real and hypothetical data science projects covering different policy areas and different data science approaches – a summary of these case studies can be found in chapter two. What followed were rich discussions about the impact that different approaches to data science have on a project's overall acceptability and the practical and ethical considerations that arise.

It is worth noting that the research was not designed to give definitive yes/no feedback on specific projects. **Case studies served as a useful starting point for deliberation** and revealed the **types of reasoning and assumptions** that lie behind people's judgements and the broad principles that they applied when assessing the future of data science in government.⁸ Similarly, the purpose of the **conjoint exercise** in the online survey⁹ was to establish the **underlying factors** in decision making, and to identify the components of projects that respondents were most positive or concerned about.

Summary

The general public were largely open to government conducting data science projects, especially where a clear wider public benefit could be established and where the risk of taking no action was significant. However, support for data science was conditional, based on an assessment of the overall value of the data science project, and whether the benefits outweighed the methodological risks.

Before fully engaging in discussion about specific methodology, opportunities for data science are needed to first pass a broader value assessment: Do I support the overall policy objective? Do I support the intended outcome of the project? Is there value in using a data science approach over or alongside other methods? When these conditions were not satisfied, or were contentious, many were inclined to dismiss projects outright, before any real consideration of how data science could make an ethical and robust contribution.

Opportunities that passed this first stage were then subject to a nuanced risk assessment of the entire project. This assessment balanced three further considerations: Is there a privacy concern? Is

⁸ The case study stimulus that was used, while carefully informed by people working in government, was not a statement of government policy that the public were being asked to comment on.

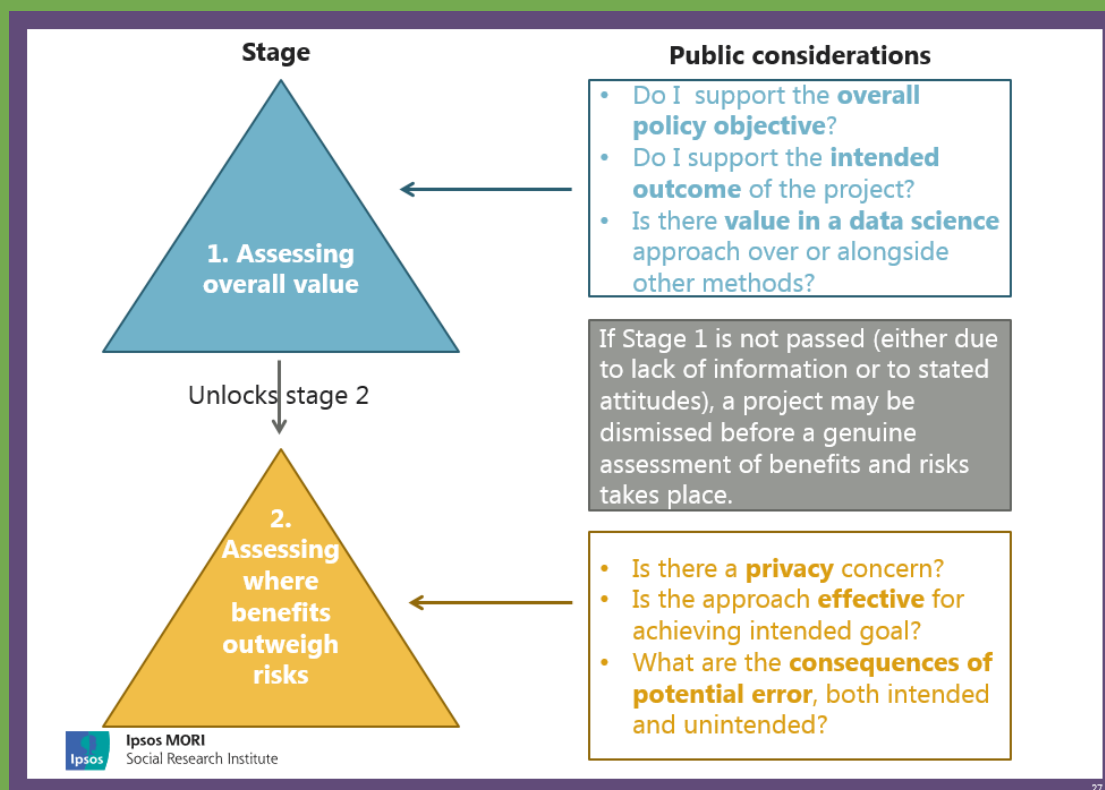
⁹ For more details on how the conjoint exercise works, please see section 2.3.

the approach efficacious for achieving the intended policy goal? What are the consequences of potential error, either intended or unintended?

Within this, concerns about risk, consequence and efficacy are measured against the specifics of the policy aim to judge the merits of a data science project. These formed the general assessment of projects, within which personal values and viewpoints further honed opinions.

Personal and ideological concerns can affect how individuals weigh up the benefits and risks of a project's features. For example, people who are more positive towards the idea of using data science in general and persuaded of the benefits for the wider public tend to be more open to trading-off their individual privacy. While those with ideological concerns about living in an increasingly surveillance society err on the side of caution and hesitate to accept a project which they believe will contribute to this issue.

These concerns were weighed against personal and ideological values in the belief in data science to deliver stated policy goals. This included views about technology and trust in technology to deliver important policy goals with less human interaction; how data science projects would interact with existing policy interventions (and especial concern if data science was to *replace* traditional methods); and debates around the need to utilise technology to keep up, either with other organisations, or other governments.



5.1 Assessing overall value of data science project

Recent research for the Wellcome Trust on public attitudes towards commercial access to health data found that a *“strong case for public benefit is the most important factor for many people: without it, data by any organisation is rarely acceptable.”*¹⁰ The Wellcome Trust work showed that the first test that people apply when assessing the acceptability of a health data-sharing scenario is the purpose of the activity: WHY it is being done. In order to pass this test, the purpose must involve some significant (if not solely) public benefit. Any data-sharing activity that results in private interest alone and displays no obvious public benefit is deemed unacceptable by most people.

All of the case studies presented during this research already contained some form of public benefit, and thus the public benefit test identified during the Wellcome Trust was easier to pass. However, simply stating a public benefit did not guarantee support for data science; opportunities were identified based on relative public benefit, and through an assessment of the following three considerations:

1. Do I support the overall policy objective?
2. Do I support the intended outcome of the project?
3. Is there value in using a data science approach over or alongside traditional methods?

As part of the deliberative process, participants were asked to rank different data science case studies on a spectrum from least to most important for government to do.

- Case studies where the link to policy action was seen as weakest tended to be classed as ‘less important’.

For example, a hypothetical case study about **understanding sexual preferences** received less support from dialogue participants because many lacked an awareness of how the government might use this information to develop a more accurate picture of the population, to monitor trends and review whether existing legislation is relevant or needs updating. Initial reactions of unease often overshadowed deeper considerations about the different data science approaches that might be involved – for example whether the project would look at identifiable Twitter data or aggregated data collected on equality forms by commercial companies.

Similarly, a hypothetical case study about understanding **future school age population** was seen by some as being less important, because there was no clear sense of how the project might help government plan school services better. For those who *did* grasp how the project could lead to an improved public service and benefit future generations, the example was considered very important.

¹⁰ Ipsos MORI report prepared for the Wellcome Trust, March 2016, <http://www.wellcome.ac.uk/About-us/Policy/Spotlight-issues/Personal-information/Public-engagement/index.htm>

- Case studies where the public questioned the need for a new approach – where they believed the status quo was working sufficiently well – where plotted in the **middle of the spectrum**. This often reflected a level of uncertainty about the importance of the policy goal and/or the need for action beyond traditional methods.

For example, hypothetical case studies about data science being used **to identify deliberate fare dodgers** or **identify motorists who break the speed limit**, elicited more ambivalent responses. Some questioned whether it was government's role to tackle fare dodging or whether it was rightly the domain of private transport companies. Others felt that these two were cases where existing approaches were good enough and they struggled to see the value a data science approach would add.

In cases where the policy goal was a more personalised or targeted service for individuals, such as **identifying individuals leading unhealthy lifestyles**, or deciding **what pension or employment support might be right for you**, reactions often hinged on attitudes towards government more generally and the extent to which people felt it was appropriate for government to intervene in their lives. Those with more personal experience of the benefits of public intervention e.g. personalised employment support from the Job Centre Plus tended to respond more warmly to these examples.

- Case studies where people could not see that anything would be lost by using data science – and indeed where they felt that we might lose something by *not* using data science – tended to be classed as '**more important**'.

At the other end of the spectrum where hypothetical cases studies where participants were simply supportive of *action by default*. **Identifying people living in UK illegally** and **identifying terrorists** were among the hypothetical case studies that tended to gain a more immediate support for data science, with the use of data science almost becoming a proxy for showing that government was doing its job and taking the matter of public protection seriously. Similarly, examples that sought to predict future levels of cancer were also cited as more important because participants felt that any contribution to this cause would be valuable.

The more valuable the policy outcome was considered to be, the more accepted the use of data science. Accordingly, not only did overall support for using data science grow when the policy area *was* well supported, but so too did support for data science methods which might involve some level of public risk. Support was greatest where participants agreed with the **overall direction of the policy**, with or without data science taking place. For example, this might be the public benefit of understanding how people are moving across the country, between censuses, or improving transport services.

Secondly, participants considered **how a data science project could contribute to that overall direction** and whether the outcome would lead to an intervention they could support. They might well agree with the policy direction, but not with the outcome of the data science project. For instance, some participants agreed that there were opportunities to use data science to help people lead healthier lifestyles, but only where this was

used to improve and target services, not to identify individuals to target messages. Similarly, some agreed with the purposes of countering terrorism, but disagreed with scenarios that might lead to people being arrested based on the evidence of data science alone.

Consideration 1: Do I support the overall policy objective?

Where participants could not see the relevance of using data science for achieving a public benefit – in other words they did not see the point of the activity for government policy-making or delivery – they were likely to voice resistance.

I don't see why they need to know if you're healthy or not, that's your life, your business.

Taunton, event 2

Participants were much more wary of giving their support for data science projects in areas where they had **limited understanding of how government makes policy and delivers services**. For example, using data science to understand the future school age population or to understand sexual preferences of the UK population, were two cases that many initially struggled to accept as they did not see how government would use the insights for any tangible public benefit. Views often shifted, however, when it was explained that this could help government to plan for future generations and build schools, or reform laws so that they respected the different lifestyle choices of people living in the UK. Explaining the relevance of how data science can add value to popular policies is thus key to identifying opportunities.

Why do they care about our sexual preferences??

London, 'high tech' event 2

Understanding the government's remit when it comes to some aspects of society also clearly affected public perceptions of whether there is likely to be a clear public benefit. For example, some participants struggled to understand government's role when it comes to privatised transport networks and food. As the following case study shows, this sometimes hindered further engagement with the details of the project and meant that participants dismissed it outright.

Case study – identifying fare dodgers

Dialogue participants were shown an example of how data science could be used to spot deliberate fare dodgers on public transport systems, and were presented with data science approaches of varying levels of intrusiveness. Several participants felt that our transport networks are run to such a great extent by private providers that there was little merit in the government getting involved. On the other hand, some saw it as an easy-win for data science since transport is a sector where private companies are unlikely to take responsibility for policing and where the proposed data science approaches seemed to offer tangible public benefit without being too invasive or in danger of creating public harm. This highlights the extent to which public acceptance of a data science project hinges on being able to see how it could, indeed *will*, lead to public benefit.

Personal attitudes towards government also influenced participant initial assessment of the project. For example, those who generally occupied a more libertarian stance to government found it more difficult to appreciate that there might be value in government using data science to offer more personalised government services, like identifying people living unhealthy lifestyles or deciding what type of employment support was right for an individual.

Personal interaction with government services was another factor that could sway people's initial value assessment of a project. Individuals who had lots of contact with services and knew the benefits of handing over lots of information about themselves – for example when visiting the job centre – as well as the drawbacks of advisors not having enough information to provide relevant support, tended to be more able to see the potential for public benefit.

Consideration 2: Does the policy outcome lead to public benefit?

Perhaps the biggest hindrance to acceptability of a particular data science project was the extent to which it might **create harm or inconvenience to members of the public or organisations**. Where significant harm or inconvenience was suspected (even without knowing all the details of a project), the project was unlikely to pass the first test of being 'in the public interest' and was thus dismissed on those grounds alone.

For example, in the case studies below, the outcome of the project affected not only whether the project met test of 'public benefit' but also the mechanics of the potential project, and how much risk they were willing to take for both intended and unintended consequences. For those outcomes with the most at stake for a single individual or organisation (for example closing a food outlet, deportation) participants were less flexible in the type of data, methods used, and balance of false negatives to false positives.

Case study – using online reviews for targeting restaurant inspections

Participants were asked their views on whether it was ok for the FSA to collect online reviews of restaurants to identify which food outlets might have poor hygiene standards. Many participants felt that the data would not be reliable enough, giving bogus or malicious reviews the potential to ruin businesses based on poor evidence. Although they were broadly supportive of the overall aim of regulating food hygiene they were not supportive of the project **if** the outcome would be to automatically shut down a food outlet based purely on the data science project. However if the FSA were using the data to help them prioritise which establishments to inspect first, participants were generally accepting of this use of data.

Similarly, while participants did not mind data being gathered to identify illegal immigrants if it was going to be used to plan local services better, they did not want the data to be used to identify individuals who might then have their liberties away. The suggested outcome of the data science project had a clear impact on whether participants thought these approaches should be used in the first place.

There was further discomfort in the belief that data science approaches that aim to identify and tackle illegal activity will inevitably include the data of those who will have done nothing wrong. Participants often reacted negatively to this concept, which they saw as inconsistent with the existing approach that citizens remain innocent until proven guilty in a court of law. Although participants appreciated the logic behind collecting information about everyone in order to identify a small group who operate illegally, they were often uncomfortable with the required trade-off.

For example, the case study 'Who lives at your house?' (which used loyalty card data to infer how many people were living in each house to identify possible benefit fraud) was questioned by participants due to the fact it was perceived to be unfair to process the data of all customers, rather than only looking at the loyalty card data of those who were suspected of potentially committing benefit fraud. Participants perceived the approach to be overly invasive, and were concerned about the risks of innocent people being incorrectly identified. The intuitive sense of unease prompted some to identify the examples of scenarios where those who might appear guilty were innocent, and as such render the data science approach unworkable.

Is it ok to match up shopping data with tax data? To a certain degree. For people bending the rules, yes, but it's not fair on everyone else. Only target the suspected.

Taunton, event 1

This points to a potential need to further engage with 'at risk' individuals and ensure that their attitudes towards government's use of data science are accurately taken into account, given that the threat to personal privacy for them may be greater but that they might also stand to lose most if the project does not go ahead.

As shown in section 5.5, the conjoint exercise shows that **whether individuals or groups will be targeted** makes a difference to overall acceptance of a project. Individual-targeting was the most likely level to negatively affect decisions, while targeting specific groups was more likely to have a positive effect. Targeting individuals

was most likely to illicit a positive response in the healthy foods scenario. This is potentially because the public are more used to receiving communications in this context, often from supermarkets themselves.

Consideration 3: What is the value of using data science compared to traditional methods?

Throughout dialogue events, participants repeatedly asked facilitators 'why', questioning the rationale for government using data and data science in different ways. This highlights the journey that the wider public will need to be taken on for any future meaningful conversation about the potential applications of data science.

Why would government want to do this? Why would they need to look at Twitter? Why would this make any difference to what happens now?

Sheffield, event 1

Experts were able to provide valuable information to groups about the reasons they might have for accessing data and using data science in their day-to-day work. Introducing this kind of information made a noticeable difference to the group's reactions, reassuring them that government had public benefit as its aim rather than simply conducting data science for the sake of it. Until participants could see the value of a data science approach, they tended to focus on specifics relating to scope of government and invasion of privacy.

Data science proposals were often questioned and picked apart by dialogue participants where **current status quo (non-data science) approaches were seen as successful and adequate**. For instance, when presented with a case study about using social media data to understand user experiences and levels of trust in the UK courts, participants struggled to see why this approach would be any better than what we have now (e.g. exit surveys and one-to-one interview approaches with people who have used the courts). Many voiced concerns that related to the type of data being accessed by government (Twitter), the quality of that data and the principle that government would use information shared in a commercial space without users being made aware. In a case like this, participants did not go on to conduct their risk/benefit assessment because the assessment of overall value had not been passed.

They'd need to convince me that they're making a significant difference to planning and making people's lives better ... I don't know enough about it, but on the basis of the scenario, it doesn't convince me. Don't think this is enough to affect infrastructural change.

Taunton, event 1

The question of whether there is a genuine need for a specific data science approach was again often affected by people's **experience of government and personal beliefs about the level of intervention or support** it should have in their private life. In the case of data science allowing for more personalised services – to provide individuals with employment or pension advice that more accurately reflected their needs – it was predominantly participants who had prior experience of the weaknesses in these types of service that supported the data science techniques, which they felt offered something better than what was already being used. A minority of participants remained sceptical of the public benefits that data science might bring.

throughout the dialogue process, usually because they occupied a more sceptical position towards government intervention in general.

Support for specific data science projects clearly grew as participants identified (and learned) the potential for **data science to not only complement but add to what government was doing already** and the increased benefits that adopting data science would bring. For example, in the case of understanding local populations using mobile phone data, there was often agreement that the Census is too infrequent to measure some changes in the population as it is only conducted once every 10 years. Most were much more positive about the case study once they considered this factor, and just needed a bit of reassurance about the type of mobile phone data that would be used (i.e. not personal text or call data).

Many dialogue participants reacted well to the idea that data science approaches might form one part of a bigger **tool-kit** that government had at its disposal. The proposal to use data science in conjunction with other approaches and that government would continue to place value on the alternatives, helped reassure people that data science was not something that the government was trying to push, gratuitously, but that it offered one credible and often effective solution to policy problems.

A step further than identifying the need for new approaches, was recognising that not only are the approaches that government currently uses in need of additional 'help', but the scale of the problem being tackled and potential for public harm if not successfully tackled, justified **trying everything possible**. Moreover, for some participants it actually *necessitated* the use of data science approaches as they offered a very powerful potential resource that it would be irresponsible not to use.

The clearest example of this was the public's response to the 'identifying terrorists' case study. Participants could see such a significant public risk and threat to human life from doing nothing new and actively *not* using data science, that many argued the case for false positives over false negatives, and used the maxim of "better safe than sorry" to explain their case for using a data science approach that might capture innocent people through its design.

"Depends on how big the risk is - the bigger the risk, the wider you have to cast your net. If it would be disastrous to miss them, then you have to include innocent people. Better that a few innocent people are a bit cross at being stopped, than a terrorist incident - because lives are at risk".

Taunton, event 2

For some dialogue participants, acceptance of data science approaches was also influenced by a hunch that to adopt data science into government delivery would be to **keep pace with the times**. Moreover, participants were uneasy with the idea that not embracing data science at the level of government policy-making and service delivery would put the UK at a disadvantage compared to other countries like China.

"I just worry what if we don't do this. Other countries will – China! In that sense we don't really have a choice, we must."

London, 'high tech' event 2

However, others were concerned about the **domino-effect** that might be triggered should the specific project go ahead. This view was usually held by participants who were more sceptical of government's intentions in general; they were prone to pick holes in the objectives of the project and argue from a very ethical perspective about key public priorities such as freedom to not be tracked or categorised. They believed that allowing this imposition on their personal freedom might open the door to a future which was ultimately not desirable.

They need to be careful they don't take a step too far because that would be life changing for everyone. That would be going into every corner of people's lives. I don't want the government knowing where I am and how I got there and how long it took. The word 'freedom' comes into it. Freedom is not being tracked and categorised.

Sheffield, event 2

This highlights the important role that **human oversight** plays throughout a data science project and the public's preference for approaches which treat data science as an additional source of insight (e.g. for ticket inspectors) rather than being seen as the final word and potentially leading to automated decisions. It also explains some of the public's **unease surrounding correlation** in driving or making decisions. Explaining how and why data can be combined analytically and with human oversight will be important for communicating the value of a data science project (and indeed of data science more generally).

5.2 Assessing where benefits outweigh risks

For those purposes which passed the test of public benefit, participants scrutinised opportunities for data science by trading off the risks and benefits of specific proposals. The conjoint analysis conducted through the online survey identified how people go about making decisions about data science projects, the trade-offs they are willing to make, and the relative importance of project characteristics.

Conjoint analysis

Conjoint analysis was conducted in order to identify the different principles that are important to people when faced with different data science projects.

As part of the online survey, respondents were asked to imagine themselves as a part of a team in government responsible for solving problems using data science techniques. Respondents were then presented with two randomly generated data science projects different scenarios where data science could take place.

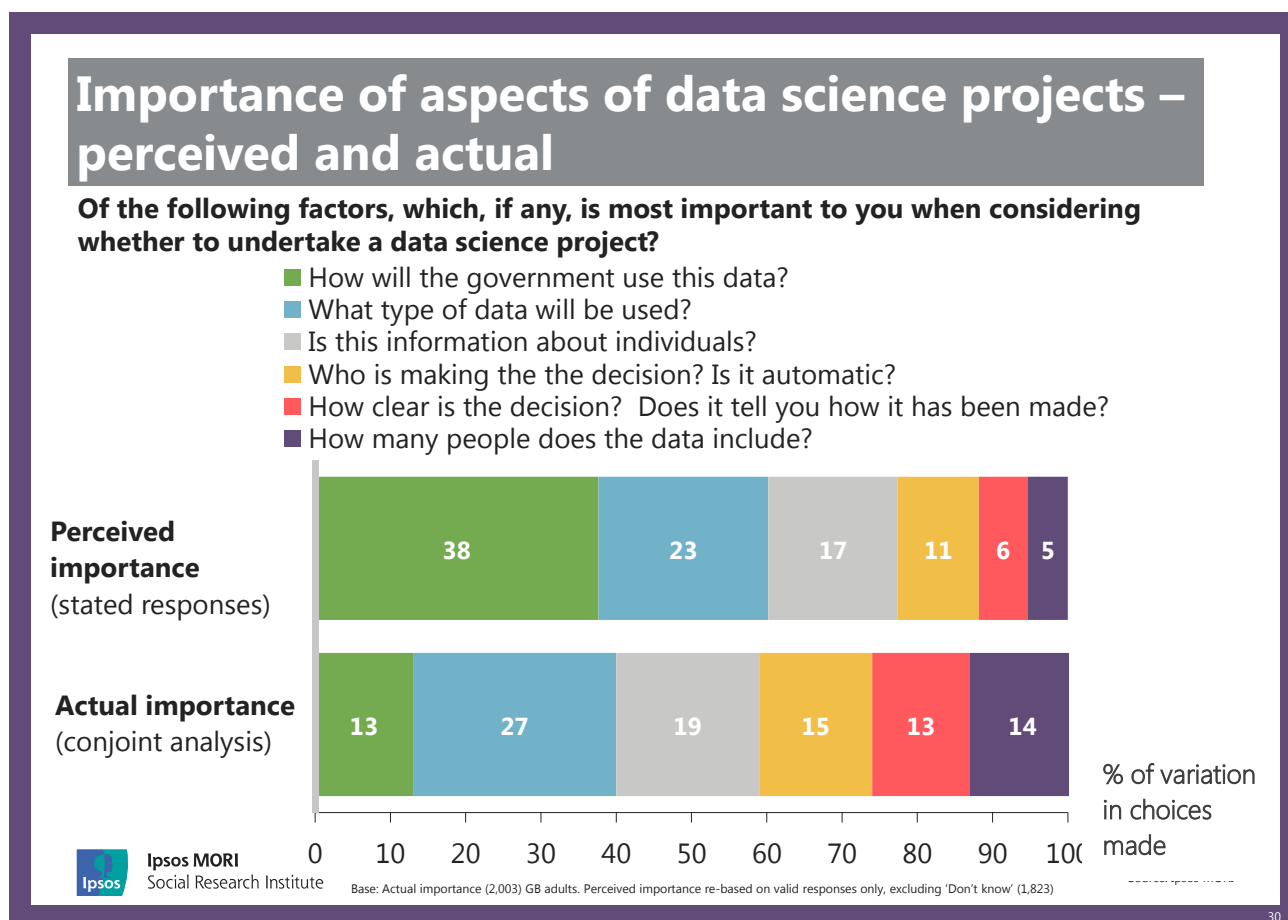
Each data science project had a description of the method that would be used across six different **attributes** (such as what data would be used, and whether the decision is automated). Each attribute had four or more **levels**, covering methods that would be more generalised but generate less specific insight vs. those which are more invasive but generate richer, more targeted, insight (such as sensitive personal information or national statistics).

A detailed description of how the conjoint analysis was designed can be found in section 2.4.

The online survey compared what respondents thought the most important factor in their decisions was ('perceived'), with the 'actual' importance of attributes as derived from the conjoint exercise.

Figure 5.1 below shows a large discrepancy between the perceived importance of 'how the government will use this data' and the actual importance derived from the conjoint. However, this discrepancy can be explained because no specific policy objective was given when asking participants what they think is important in making decisions on data science opportunities overall, and thus respondents reacted by stating that having a policy objective was paramount. In contrast, each scenario presented within the conjoint exercise had a stated policy objective, and thus other characteristics were more important in driving decisions.

This reinforces the first stage in the assessment journey explored in section 5.1, and shows the importance of establishing the policy objective in the process of making a decision. Once this policy objective has been established, as was the case in the scenarios presented through the conjoint exercise, respondents then went on to assess the other aspects of the data science project.

Figure 5.1: Perceived importance and actual importance of aspects of data science projects

Having been given a policy objective as part of each scenario, the selection of preferred data science approaches during the conjoint exercise was driven largely by the type of data that would be used – this accounted for 27% of variation in selection. This was followed by whether or not individuals could be identified in the data, which accounted for 19% of variation. A further overview of the relative strength and positive/negative association with each of these attributes and sub-levels can be found in section 5.3.

Assessment of these different characteristics of data science projects can be summarised through three considerations:

- Is there a tangible privacy concern in the project?
- Is the project efficacious for achieving the intended goal?
- What are the consequences bound up in error?

Risk assessment: decisions based on privacy concerns

Within the conjoint introduction text, we emphasised the importance of trading off privacy concerns with the effectiveness of data science projects. That is, respondents were asked to consider these trade-offs, while still allowing them to make decisions based entirely on privacy or effectiveness if that was their preference. In practice, the conjoint results show that most people weighed these two elements of projects up together.

It was possible to see that across the different policy contexts, there were clear concerns around options where there was a risk to privacy. This may also drive why **type of data** and **identifiability in the data** were the most important 'attributes' in driving people's decisions (see above).

In the conjoint analysis, it was possible to see a general trend towards options that appeared less privacy sensitive. For instance, respondents were less likely to choose data science approaches in which individuals could be identified, but were more likely to choose approached that used anonymised data or data grouped into large numbers. Similarly, more personal data, such as criminal records and sensitive personal data were chosen less frequently that approached which used data like traffic and transport use and official statistics.

Caution over the use individual level data was also present in the workshops. Many participants requested anonymisation by default, and felt uneasy about the risk of having multiple fields in a data set that could be combined to identify individuals, even when names were not present. As explored in chapter 2, participants were often unaware of how individual level data could be used either to create aggregate datasets, or to link variables and profile anonymous data for richer insight. Those who were most concerned about privacy were not reassured by guarantees that results could be generalised or aggregated. Instead, concern about data being lost or hacked led to a general unease about individual level data being collected and used by data scientists. However, as explored in chapter 6, not all people are put off by individual level data, some do see real value in the potential insight. For this group, analysis of individual level data was a positive driver.

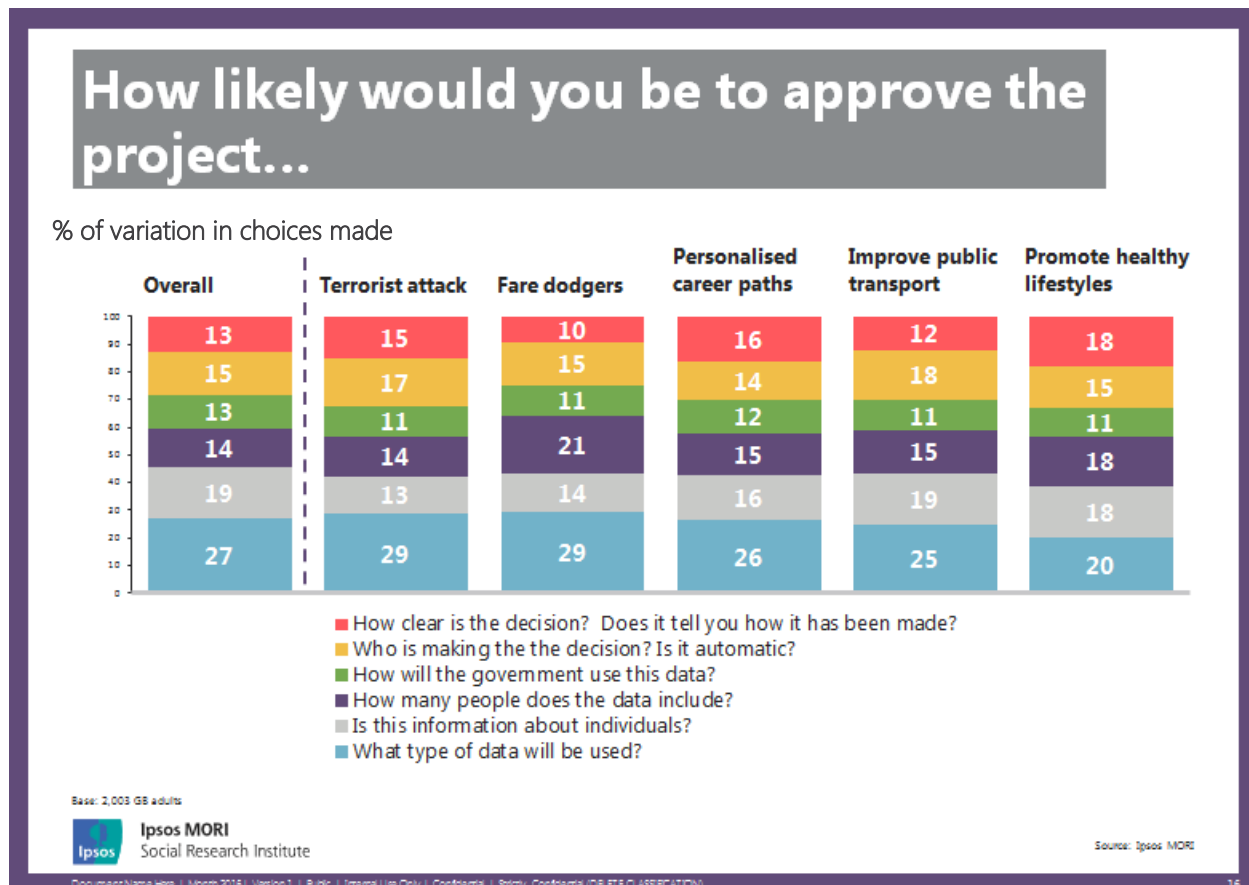
The relationship between privacy concern and the risk assessment was not entirely straightforward. While individual, anonymised data and data grouped in a large number of people both had a positive impact on decisions, respondents were less likely to choose approaches that contained data that was 'not related to people or humans'. This demonstrates that aside from simply assessing the depth of the privacy concerns, respondents were also considering other factors when making their data science decisions, such as how effective and successful the approach would be at meeting its intended outcome.

Risk assessment: decisions based on efficacy

In weighing up the benefits and risks, people also considered how effective the suggested approach would be, and were often willing to suspend some of their privacy concerns where this would add value or be a necessity to the success of the project.

Evidence for this can be seen in the below chart, which illustrates how the importance of the different attributes changed based on the policy context people were shown. For instance, in the fare dodgers scenario, respondents placed a greater importance than average on ‘the number of people included’ in the data science project. On the other hand, respondents who saw the fare dodgers scenario placed relatively less importance on ‘how clear is the decision?’

Figure 5.2: Importance of different attributes in decisions made about projects across scenarios



Digging deeper into this example, we can see further evidence of decisions being made based on the effectiveness of the options to the specific scenario. Again, for the fare dodgers example, respondents were more likely to select projects that looked at data from people in a specific area of the country. This may be because this offered a certain regional specificity that fits with people’s perceptions of trains being a local concern. Conversely, for the fare dodgers scenario, respondents were less likely to select projects that used data that was ‘not related to people’. This shows that, while this might be the most privacy conscious option, respondents felt that it was not likely to deliver the stated policy aim of catching fare dodgers.

The idea that respondents might be considering efficacy in relation to specific projects can also be seen in the counter-terrorism example. In the detail here, it is possible to see that individual-level, identifiable data was more likely to be selection than it would in the other scenarios. When counter-terrorism was the stated purpose, respondents also react less negatively towards identifiable data than they do towards aggregated data, despite aggregated data being the less privacy sensitive option. This may be because they feel the

consequences of not doing enough for counter-terrorism is much greater than the potential privacy concerns of using individual, identifiable data.

Considerations of effectiveness of data science approaches was also apparent in the workshop discussion. Here, participants focused mainly on the quality of the suggested data, the techniques used, and the conclusions that could be drawn from the data are likely to be accurate. For example, most participants were not convinced that the comments left on food discussion forums would be accurate enough to draw conclusions on the hygiene of food outlets, or that the sales of books related to cancer could indicate health levels. The effectiveness of suggested approaches was also judged against whether the data was to be used for research or operational purposes.

Risk assessment: assessing risk through error, unintended consequences and intended consequences

The third consideration applied by members of the public was assessing the scale and impact of consequences, both intended and unintended, of a given project. To help explore this further, a split sample question was asked in the online survey to compare the likelihood of approving a specific data science project when different risks are presented.

Online split sampling

As part of the online survey respondents were split into five sub-samples, each shown a slightly different version of the same hypothetical data science example. The basic example asked respondents to imagine they were part of a team in government responsible for using data science to tackle benefit fraud by combining data on an individual's tax and earnings with data on the individual's benefits. Data science, respondents were told, would be used to identify patterns in the data.

Around 400 respondents received each of the below additions of text:

- **Control group:** [no extra text shown]
- **Large number, small impact:** "Your colleague informs you that there is a risk that a large number of people may experience a small negative impact as a result of the project. For example, individuals who are innocent of committing fraud might have their data examined."
- **Small number, large impact:** "Your colleague informs you that there is a risk that a small number of people may experience a large negative impact as a result of the project. For example, a small number of individuals who have not committed fraud could have their benefit payments paused while the issue is investigated further."
- **Risk of error, with an automatic outcome:** "Your colleague informs you that there is a risk that answers will sometimes be wrong (for example, the data is 6 months out of date, and there is a chance that data could be linked incorrectly), but it will not be possible to tell until someone's benefits have been incorrectly paused."
- **Risk of error, with outcome investigated:** "Your colleague informs you that there is a risk that answers will sometimes be wrong (for example, the data is 6 months out of date, and there is a chance that data could be linked incorrectly) but that no benefits will be paused until a member of staff has investigated the issue further."

Please note, the results of this exercise are specific to the scenario presented and may not necessarily reflect overall attitude to assessing risk for all data science projects.

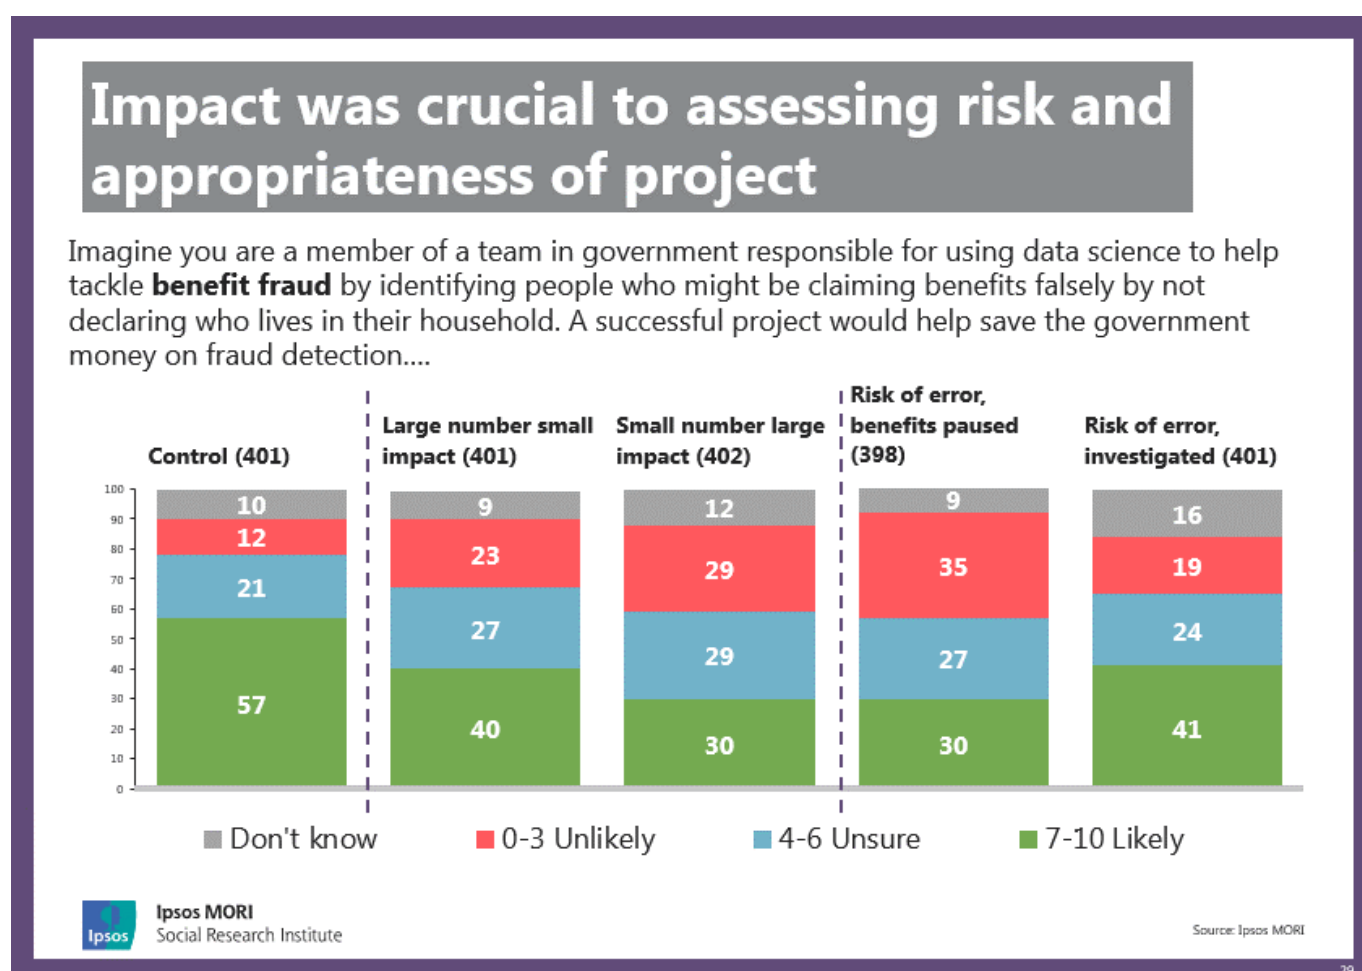
As shown in the below chart, the **control group** are more likely than not to approve this specific data science project (57% rated this project as a rating of 7-10 likely to approve). There is, however, a drop between the control and all the forms of risk we presented different sub-samples of respondents with. This suggests that risks are not top of mind for most people, and need to be stated clearly when identifying opportunities for data science.

In this context of this specific example, people were more likely to approve a project where there was a risk of a small impact on a large number of people (40%) than a project where there was a risk of a large impact on a small number of people (30%). In the example of fraud being investigated, this means that a small loss of privacy for a larger group of people was preferable to a small number having benefits paused. It should be

noted that *both* of these sub-samples were less likely to approve these projects than the **control group** sub-sample where risks were not specified.

The other kind of risk we investigated was where the risk of incorrect data was specified. However, the two groups saw two different versions of this: one where there is a negative impact based on an automatic decision to pause benefits (**risk of error, with an automatic outcome**), the other group saw an example where the decision to pause benefits is investigated before any actions occurs (**risk of error, with outcome investigated**). Again, this distinction appears to impact on people's likelihood to approve this project, with an investigated decision more approved (41%) than an automatic decision (30%). Together with the first two examples, this shows that there was a greater willingness to approve data science projects that contain human oversight over those which have automated outcomes.

Figure 5.3: Split sample differences when different kinds of risks presented to respondents



Looking at consequences was a key part of the discussions in the workshops. Privacy and efficacy concerns were seen as interrelated and people tended to weigh them up in tandem. In some instances, privacy concerns and efficacy concerns were not traded-off but were very much the same thing. This was often where there were concerns about discrimination in the project. This violated both sensibilities around privacy, but also

raised concerns that if that kind of ulterior motive was to affect the outcome of the project, the purpose of the project would not be achieved.

This could demonise certain groups, if a certain type has an ethnic cuisine then that could cause problems in communities, if the FSA are just closing down restaurants [without human investigation first]

Sheffield, event 1

Indeed, there was also some concern that if behaviour is being monitored it could change our behaviours and then modify the risk assessments of others who might be wary of such projects taking place. This argument was often raised against government surveillance as well. If you know you're being watched, you might change your behaviour, which in turn affects the efficacy of the surveillance.

I'm not convinced people will be inclined to post a comment if it could be used by a government department. I would think it's a step too far and I wouldn't want to be involved as a reviewer in having something shut down. I would hold back.

Sheffield, event 1

Participants generally found it easy to start weighing up the risks of unintended errors taking place. This was often, but not always, framed in terms of the error that could occur if human oversight was not in place. While some of this concern was a broad fear about new technologies, there was also a more nuanced discussion of why decisions about consequences needed to be made when using technology to implement a policy goal.

I had an experience where I set up a direct debit, I came home to 200 letters on my mat, there'd obviously been a mistake with the computer, and why wasn't someone at the post office saying there's something not right?

Sheffield, event 2

Other concerns about error flowed from the risks of datasets being inadequately secure against, for example re-identification. This depended on the standards that should be demanded depending on the intrusiveness of a project, and the efficacy of the project to deliver. For example, security was more important when there was a risk of re-identifying personal information, but support was reduced for using de-identified data if it reduced the overall efficacy of the project.

Like when [EXPERT] was saying with the mosaic, you need to consider the risk of identification. It doesn't mean don't do it, just realise at the end of the day that the data you are producing could have implications.

Taunton, event 2

The principles of intended error, also characterised in terms of 'false positive' and 'false negatives', needed to be introduced to participants. Once they were introduced these forms of intended error became an interesting way of discussing the trade-offs that are necessary when building a data science project. In general, the

concepts of false positives and negatives resonated with participants, who were able to enter into some of the terminology and metaphor that often is used to simplify complex technological concepts.

There was no golden-rule established during the workshops; instead, participants consciously realised that their willingness to accept intended consequences was directly related to the importance of the purpose of the project and to the seriousness of the consequences for those who would be wrongly identified.

Most participants preferred an approach that would have a small effect on a large number of people when identifying people on a no-fly list – in this example the risk of not capturing all people on the list was greater than the consequence of delays incurred by innocent people wrongly accused. A preference for more false positives was also expressed when looking to improve careers advice – here participants preferred to see a higher number of career ‘matches’ which they could filter themselves rather than risk being shown only a small number of matches that highly correlate to their characteristics.

Depends on how big the risk is - the bigger the risk, the wider you have to cast your net. If it would be disastrous to miss them, then you have to include innocent people. Better that a few innocent people are a bit cross at being stopped, than a terrorist incident - because lives are at risk

Taunton, event 2

However, while a lot of the discussion came down on the side of producing more false positives, there were also scenarios where people wanted to err on the side of producing more false negatives. This included when the risk to the individuals who might be wrongly included within the ‘net’ was greater. This was often the case where the outcome for an individual was great, but the outcome to society in taking action against that individual was relatively small – for example, identifying benefit frauds or illegal immigrants. Here participants needed greater confidence that individuals would not be wrongly accused.

There's a risk of alienating people and pushing them towards being extreme. It will cause resentment and anger"

Taunton, event 2

Personally I think false negatives would make more sense because you're getting people who definitely did it and not bothering anyone else, whereas false positives is more important for serious offences because the consequences if someone slipped through the net are more severe.

Taunton, event 2

5.3 Exploring the features of data science projects

As noted above, once the core purpose of the policy or data science project had been accepted, a more sophisticated discussion around the features of an individual data science project could take place. In this second stage, where the benefits and risks of individual proposals are considered, small changes in the features of a data science project could alter the overall assessment of acceptability that people make. Although there

were no general red lines for assessing features across data science projects in general, some features had a greater baseline of acceptability than others; furthermore, all features were considered in the context of the original policy objectives and outcomes. Thus, more distinct 'red lines' did appear within specific individual scenarios of data science projects; changes in features could turn what might otherwise have been an acceptable project into one that raises concern, and likewise push one that might be problematic towards a greater level of acceptability.

The table below explores both the general baseline of acceptability for different features of data science projects, and provides examples of how small changes can make large differences in whether an approach is perceived as acceptable and effective. The table should be used as a helpful starting point for the evaluation of future data science opportunities, but not as a sole reference point. There are no hard and fast rules driving public acceptance, with public support tied to the context of the specific policy objectives and outcomes of individual projects. As such, it is possible that a project that contains characteristics with a lower baseline of acceptability would be seen as acceptable where the policy is supported and the approach deemed effective; and vice versa, where a project that contains characteristics with a higher baseline of acceptability would be seen as not acceptable where the policy is not supported or the approach is deemed ineffective or unnecessary. Furthermore, the distinctions between characteristics are **relative**, and thus one having a greater baseline of acceptability than another does not necessarily equate to public 'support' for the concept.

	Choices made through conjoint exercise	Feedback from workshops
Type of data used	<ul style="list-style-type: none"> • Respondents were most comfortable towards use of public non-personal data (such as level of traffic, use of transport, job vacancies) and aggregated and anonymous official statistics. • And least favourable to use of sensitive data such as race or ethnicity; however, this was more acceptable when used to personalise services rather than target individuals for wrong-doing. • Respondents could see the value of using personal transactional data (for example, travel card or store card data) and other personal data (for example criminal records or phone location) for some scenarios, but only where the data was highly relevant to the project objective, such as improving public transport or preventing a terrorist attack. • Overall, use of social media data was a negative driver when selecting potential data science projects. 	<ul style="list-style-type: none"> • Non-personal, non-sensitive data (e.g. car park sensor data) was widely accepted, although the fact that it is collected into a dataset was not top-of-mind for many. Some were sceptical towards this kind of data as they question whether it really is non-personal – e.g. when shown example of car park sensors, some jumped to the conclusion that it would pick up information like registration number – while a misperception this indicates where there is some paranoia (and highlights need to underline when data is non-personal/non-sensitive for public trust). • Discussion of personal transactional data usually led to questions about why government would want access to this kind of commercial data (e.g. store cards); this represents a ‘context collapse’¹¹ as the public are not expecting government to enter this space and many do not see the potential value that the data could have. Some worry that it reflects a ‘surveillance state’. • Social media data often elicited ambivalent responses – some participants had an intuitive sense that social media was not government’s space and that users are not made aware when they share information on there that it could be used for this purpose (again, this represents context collapse). Others saw it as unproblematic since social media forums are inherently public spaces so users shouldn’t have false expectations of privacy. All social media data was often clubbed together as

¹¹ See p40 of Ipsos MORI report for Wellcome Trust on Public Attitudes to Commercial Access to Health Data
http://www.wellcome.ac.uk/stellent/groups/corporatesite/@msh_grants/documents/web_document/wtp060244.pdf

the same thing e.g. personal profile on Facebook, all tweets with the word 'crown court'. Again, while this is a confused view it shows the sensitivities around social media data that exist for some individuals. The use of 'aggregated' social media data e.g. number of tweets made in Glasgow between 5-8am, while not seen as sensitive raised a question of relevance 'Why would government need that? What can they possibly do?' and where a question of relevance emerges so too can worries about over surveillance. However, participants were impressed by the predictive power of the FSA's project identifying risk of norovirus epidemic, this emphasises the need for the value of using social media to be clearly specified.

- **Sensitive or 'private' data** about an individual's lifestyle choice (e.g. sexual preference) or state of health (e.g. unhealthy lifestyle) was also a cause for concern as it indicated that government was being 'nosy' and/or potentially making assumptions about a person without knowing the context. Here the worry around relevance leads to a worry about what government might do with certain information and how it could negatively affect an individual or group of individuals. Most participants found it difficult to comprehend how sensitive or private data could be used in a safer environment (such as aggregated or anonymised).

Level of information used

- Overall, respondents were comfortable for **individual level data** to be used as long as people **could not be identified**. This was perceived to be just as valuable as aggregated data based on groups of people.
- Most people were cautious about use of data where

- **Aggregate, non-identifiable data** largely accepted so long as the users of that data could be trusted not to attempt to identify someone by linking up different datasets. In certain contexts, e.g. healthcare, many expect that this kind of data is already being used by government so using it for data science purposes is unproblematic.

individuals could be identified, however this was more acceptable where there was an obvious operational need.

- **Individual and identifiable data** was more accepted by participants with experience of how government having access to this kind of personal information could be beneficial for an individual. E.g. in the case of providing more personalised employment advice, participants who knew how valuable it would be for an advisor to be able to link up their employment history with others' were generally more supportive. However they were more cautious towards using individual and identifiable level data where the intervention was more beneficial to society rather than the individual (for example identifying fare dodgers).

Scope of dataset – the size or the coverage of the data

- There was less discrimination in the **scope of datasets**, driven largely by considerations of effectiveness rather than ethical concerns about consent.
- For example, use of data focused on specific areas of the UK was a strong positive driver for identifying fare dodgers. Moreover, use of data that was not related to people (such as operational data, food sales, train passenger numbers) was associated positively with promoting healthy lifestyles, but negatively with identifying fare dodgers.
- **Consent was not necessarily a prerequisite** for use of data, and in some cases was perceived as a negative driver – for example in preventing a terrorist attack.

- **Size of dataset** was not a pertinent point for participants at dialogue workshops.
- **Impact of different consent models (e.g. opt-in/opt-out)** on the scope of a dataset was not spontaneously picked up. Nor did participants deliberate much over the level of insight that a dataset could therefore provide and the impact that might have on how effective a data science project would be at achieving its policy goal. This could be because it is relatively less important than other issues such as privacy, but also because it is a more difficult concept to understand and trade off. This does not suggest that it is not at all important.

Purpose of using data

- Respondents were most comfortable with using data science to **target specific groups** of people, or to **research** groups and areas *in need*.

Support for purposes with clearest, unproblematic link to public benefit.

- If policy goal provided a **better service for individuals** e.g.

- Most were less comfortable with using data science **to target individuals** for contact. This was more acceptable in some scenarios where there was clear benefit to an individual or society (such as promoting healthy lifestyles or preventing terrorist attacks), but this principle was not universal and associated negatively with identifying fare dodgers and personalising career paths.

personalised employment advice or pension support, participants tended to be more likely to give support.

- Likewise, where the goal was clearly **wider public protection** participants were more likely to be favourable. For example, in the ONS case study where mobile location data is used to understand population movements, people would reassure themselves that the project was valuable because, for example, it could help to better inform police resource and thus they could directly relate to the potential benefits.

Concern over purposes where link to public benefit is less clear e.g. research.

- While projects with a research purpose often posed less of a tangible risk to individuals e.g. understanding sexual preferences of the population, participants often applied caution as they did not see a clear rationale for government's actions, which could lead them to fear the worst (either that the UK was becoming a 'nanny state' or, for a few participants, that individuals might be discriminated against).

Relationship between human and computer

- In all scenarios, respondents were least comfortable where **decisions are automated** with no human oversight or control.
- However there was little differentiation in the **role of the human**. Having some human oversight to intervene if things go wrong was favoured just as much as humans using data to make their own recommendations and decisions.

- **Broad consensus that human oversight is essential.** A core public principle of conducting data science work *at all* was that the process should be humanised.
- Few participants had clear conception of how that would look but it was clear that they wanted an element of **human involvement at each stage** of a project including designing algorithms and monitoring them to deciding what action to take based on the project's outcomes.
- Many **intuitive fears around leaving too much to**

automation and allowing the computer to decide without a human-check.

**Clarity of
computer
decisions**

- Overall, respondents were more comfortable with data science projects where **workings of the computer were known to staff** using the data, or where they were fully transparent. Knowledge of the creators alone was not seen as sufficient.
- However, full **transparency was not a necessity**. With the exception of personalised careers advice, data science projects where those working with the computer know how it learns and makes decisions was perceived to be just a positive as a fully transparent system.

Participants found it difficult to comprehend the mechanics of how computers make decisions during data science, and the extent to which these may not be transparent or easy to determine. Despite this, there was a clear preference for ensuring that data science projects were accountable. Participants were more supportive of data science where errors could be challenges and the decision process fully understood. They expected relevant checks to be in place to minimise unintended errors and consequences.

As noted above, certain features of data science projects clearly have a higher baseline of acceptability, while others are more likely to prompt resistance. However, it is also true that small changes to an individual project scenario can have a large impact on whether an approach be perceived as acceptable and effective. Some distinct red and green flags appear within project scenarios, depending on specifics such as 'type of data used', 'level of information needed', 'scope of dataset', 'purpose of using data' and 'relationship between human and computer'. By making changes one of these features, public support for a project can wane, or it can grow, turning what might otherwise have been thought of as a problematic project, into a much more acceptable one.

The following four examples demonstrate how relatively small changes to the project's features made a big difference to levels of acceptability recorded during the workshops discussions.¹²

1. To decide what type of employment support is best for you

Public acceptance was usually high at the start, due to widespread support for the overall policy objective of providing individuals with tailored careers guidance. However, support was not unconditional, and changed based on application of false positives/negatives:

- Where social media (LinkedIn) data was used to segment populations and target individuals with one specific career pathway, according to what 'group' their characteristics suggested they fitted into, people quickly withdrew support due to lack of trust in the information presented – they did not want to be segmented.
- Where individuals were offered several career pathways based on a less rigid algorithm, openness to using social media and historic government data to make predictions often grew. In this instance, people were happier with an approach that allowed for false positives. By including options that might *not* be suitable for an individual, people often felt that they were being given greater autonomy in the decision and they felt this guarded against individuals being wrongly classified through segmentation.

2. To help tackle benefit fraud.

Most participants were cautious of this project at the outset because they felt it was not appropriate to use social media data to identify benefit fraud. However, support increased without the use of social media data.

- The use of non-government held, social media, data to identify and target individuals who was strongly disliked due to the perceived infringement on individual privacy, but also because they felt it would be an inaccurate source of data, with a high risk of error attached to an important outcome.

¹² Further information about each case study can be found in section 2.3

- However, participants were more supportive of the suggestion that the project could be conducted by linking government datasets together without social media, and where human oversight was used to investigate and sense check results prior to taking action against individuals.

3. To identify motorists who break the speed limit

Public acceptance for this project tended to begin relatively high, due to its clear, tangible link to the public benefit of saving lives otherwise lost through speeding. But there was a noticeable shift in comfort levels once the project was seen to infringe too much on personal privacy:

- The idea of placing smart sensors in vehicles to constantly monitor an individual driver's speed, struck people as a step too far. Whilst it was recognised that this feature would increase the effectiveness of the project, it was widely dismissed as being overly invasive and support for the project disappeared.

4. To predict future levels of cancer across UK

Much like the examples above, initial responses to this project were positive, again due to its obvious link to public interest and association with public good. The type of data used in different scenarios however elicited very different responses, with commercial and personal transactional data raising concern:

- The use of commercial data for online sales of books with a health focus puzzled many participants who questioned the relevance of this to cancer levels and were unconvinced of the project's effectiveness. Like other sources of online commercial data, such as LinkedIn, this space can intuitively feel like non-government territory as it represents 'context collapse' for the public who are not expecting their online purchasing habits to be subject to government scrutiny. While this is a false assumption given government would not, in this scenario, have access to individual level transactional data, it reflects the levels of ambivalence surrounding government's use of this type of data.
- However, the use of aggregated supermarket data was preferred to personal supermarket data, as it removed the risk of identifiability. But even in aggregate form, supermarket data was still 'transactional data' which people were not expecting government to access and many struggled to see the relevance of food shopping habits to cancer levels.

These examples demonstrate that there are no hard and fast rules driving public acceptance of data science projects. It is therefore important that decisions about specific features of data science projects are not made in isolation.

5.4 Implications for data scientists and policymakers

A level of scepticism was apparent across almost all participants, permissive or otherwise, when the purpose of a data science project appears to lack tangible public benefit. This points to the importance of clearly **communicating the intended public benefit when running a data science project** and being aware that members of the public are looking for this condition to be met before they will give their acceptance.

Furthermore, the public clearly view **data science within a context of wider policy decisions** and were unwilling to carry out their own risk assessment without knowing some facts about what kind of decision government could make based on the data science approach and what actions might be taken. Thus, data scientists should be mindful that they are not operating in a vacuum and strive to collaborate with those who work closer to policy to discuss what the implications are and discuss opportunities that data science approaches allow.

Participants' responses to different projects also revealed very varied levels of understanding about **how machines and humans interact**. Without good communication of this interplay and the points of responsibility that are still held by a human, there is a real risk that people will expect the worse (for example, machines making decisions automatically without humans questioning the data or data science throughout). Dialogue participants made it clear that in order to support data science projects, they needed to know that there was a "*human face*" to data science and that a layer of human oversight was built in from beginning to end.

Overall, given the assumption among many that government is working with public benefit as its aim and accordingly the fact that **red lines may not be the most helpful way to conceptualise how to protect public interest**. Nevertheless, the public do apply a **broad set of principles** to the whole discussion and thus the framework approach to regulation seems well-suited.

The final chapter reviews the public's priorities for an ethical framework in more detail and compares and contrasts what participants discussed spontaneously with their reactions to a pared down version of the GDS draft ethical framework.

6. Differing perspectives

6 Differing perspectives

Though the stages of decision making outlined in chapter 5 were common, not everyone reached the same conclusions when judging specific data science project scenarios. This chapter explores the different types of personas and perspectives witnessed during the dialogue.

Summary

Public perspectives on government using data science differ noticeably across individuals. Future engagement and communications with the public on the issue need to bear a number of factors and perspectives in mind, in particular the **impact of someone's lifestyle, past experience and political philosophy** on how they receive data science in government.

Differences do not align neatly to age and as other work on data shows the relation between age and openness to Big Data and data-sharing possibilities is complex, with factors such as level of interaction with government services or ideological conceptions of the relationship between citizen and state carrying more weight.

A segmentation analysis of the conjoint exercise revealed four groups, with variations in attitudes and their likely responses to distinct features of data science projects. There are: **'data adopters'** who support using data science for research purposes and see how individual level data can generate better insight; **'data adapters'** who respond best to uses which improve services for individuals and use non-sensitive data; **'data pragmatists'** who are more ambivalent in their views, wanting government to explore new ways of using data but are most comfortable using data for high-level statistics rather than advanced data science; and those who are **'data wary'**, who apply caution to the principle of data science and are liable to reject projects based on concerns of privacy, effectiveness or limited understanding.

Workshop discussions further suggest that **experience of sharing data** either with government or the commercial world has the potential to drive support for government data science projects. But experience of sharing data still needs to be accompanied by **exposure to the personal or public benefits** that a project might bring – hence why individuals who are direct recipients of public services often give more support.

6.1 Public perspectives

Analysis of the survey and workshops reveals that public perspectives on government use of data science are shaped by a number of different attitudes and experiences:

- attitudes towards government and policy,
- contact with government services,
- awareness of data science,
- attitudes towards the value of data science overall,
- the perceived baseline level of acceptability in specific feature of data science projects, and
- the perceived relative importance of different features in comparing risks and benefits.

Differences in view transcend age, as other factors – such as contact with government services or overall attitudes towards government – tend to supersede the effect that age might have and carry greater weight in determining how accepting someone will be of a data science project.

A segmentation of the results from the online survey identified four broad groups based on the attitudes and experiences that also emerged in the workshops: 'data adopters', 'data adapters', 'data pragmatists', and those who are 'data wary'.

Segmentation

A segmentation analysis was conducted based on the responses of 2,003 adults aged 18+ to the online survey. The analysis compared:

- the relative importance individuals gave each of the **attributes** presented in the conjoint (for example 'type of data', 'what the government would use the data for'),
- the extent to which individual **levels** (for example sensitive personal data, personal transactional-data were either positive or negative drivers in their selection of data science projects, and
- whether individuals were comfortable sharing personal information to receive personalised services.

The segmentation identified 4 clusters of respondent. These clusters were profiled further by comparing their responses to other relevant questions from the online survey.

Data adopters accounted for 23% of those completing the online survey, and were the most positive about potential use of data science by government. Although the 'type of data' used in data science projects was still their top driver when reviewing opportunities for data science, this was less prominent than in other segments. In addition, data adopters were the most interested of all groups to carefully examine the mechanics of how data science projects will work – particularly the extent to which decisions are computer or human led – and how they will achieve their objectives.

Data adopters were particularly supportive of utilising data for research and recognised the value of individual level data to generate this insight. However, they were less comfortable with projects that would use individual level data to target individuals for further contact.

Data adapters accounted for 28% of respondents, and could also see value in using data science. This was especially the case for projects that could better target individuals or personalise services, and where the data used would be volunteered by individuals or already relatively public. Although they favour opportunities that are more operational than research based, they are very cautious about the use of more sensitive personal information, and other personal transaction data (such as store or travel cards); this drives many of their decisions about opportunities for data science.

Data pragmatists accounted for 27% of respondents. They could see the benefits of using data for high-level statistics, but felt that individuals should not be identifiable in the data unless there is a clear need. Like data adopters and data adapters, data pragmatists wanted government to search for new ways of using data; however, in line with concern about use of identifiable data, they were more reluctant to share their own data for personalising of services. They felt reassured that government take these concerns seriously.

The final 22% of respondents were **data wary**. People in this group were the least likely to identify opportunities for data science and wanted more information before they made a decision about a data science project. Those within this group who did identify opportunities for data science, were cautious of the use of any kind of individual-level data in data science, and instead looked for opportunities where aggregated data or non-human data could be used.

During the conjoint exercise, around a quarter of those who were data wary chose not to undertake any form of data science project for all four of the scenarios they were shown. Among this group, the regular rejection of data science projects can be explained through three different perspectives:

- Privacy concern: some were concerned about the need to undertake any form of data science, often referring to use of data by government as 'Big Brother' and citing concern with perceived 'invasion of privacy'. Some of this privacy concern was linked to distrust of government, but also a more general fear of giving out any personal information.

I do not wish my details being given out

Response to online survey open ended question

I don't trust the people into whose hands it will fall

Response to online survey open ended question

- Effectiveness concern: others in this group were concerned that the data science projects would not lead to the stated objective, and commented that the options presented to them would have limited applicability or give misleading results. For many in this group, it was difficult to understand and believe in the value and capabilities of data, and data science techniques.

It's not that I do not like the idea of data science projects ... just could not agree with all the criteria in each option so would rather not make the decision.

Response to online survey open ended question

The options don't seem to be able to give the data needed to give a true answer.

Response to online survey open ended question

- Lack of understanding or information: For some, there was a difficulty in understanding the concept of data science, others asked for more information to help them make an informed decision. For this group, the decision not to take any action was a genuine 'don't know' option. Those who opted not to undertake any form of data science were most likely to say they have 'never heard of' data science as a term (55% compared with 35% of those who did not reject any data science), and most likely to say that they 'don't know' either whether government should make better use of data or whether they would be happy to share their own data in return for personalised services.

I need far more info before any preference decision could be made.

Response to online survey open ended question

6.2 Does experience of sharing data breed favourability?

At the outset of the project, there were two hypotheses relating to what factors might play a part in shaping public perspectives: **current data-sharing activity with government** and **digital media literacy**. Individuals who share lots of information about themselves with government might be better able to foresee the benefits of a new approach and be more aware of existing limitations. While those more fluent in digital media and technology might have greater appreciation of the opportunities offered by data science.

Sampling and recruitment

To explore the relevance of current data-sharing activity on views, a group of 'High data interactors' was recruited and screened according to the number and regularity of interactions they had with government services [this was used as a proxy to identify those more likely to be sharing personal information and more familiar with receiving a public service in return].

To explore the relevance of digital media literacy, a group of participants were recruited in London that was screened on usage of digital media and technology. (See Appendix for recruitment screener.)

High data interactors

Members of this group were noticeably more relaxed about the overall concept of government using data science and accepting of projects where this included government accessing their personal and identifiable information. This was particularly apparent among younger members of the group and those who had more contact with services that involved a level of personalisation (e.g. Jobcentre Plus).

Some *high data interactors* were relatively warm to the idea that data science could be used to make judgements about what kind of service suited them best. They had fewer qualms with government using historic data to identify different types of service user and make inferences.

For many in this group, data science done responsibly represented an efficient and joined-up government that is keeping pace with the private sector. Indeed, several felt that government was being overly cautious in consulting the public about the issue in light of the scale of data-sharing that takes place in today's commercial world. They also had an expectation that government *should* be doing this sort of work already.

"They should be talking to each other, the right arm needs to know what the left arm is doing".

Wolverhampton, high data event 1

"I don't think the government should be that bothered about it, because we are opening ourselves up to it all day every day ... I think the government is being very polite!"

Wolverhampton, high data event 1

High tech, high digital

Some members of this group were quicker to grasp the abstract value of data science, being more familiar with Big Data and how it is widely used by commercial companies like Amazon and Google. But despite this, most still needed support to understand the full scope of data science, what it could lead to and what practical and

ethical considerations were involved. Thus the process of education and contextualisation outlined in chapter 4 retained value for this group.

There were individuals across other workshops who presented themselves as more technologically and digitally literate, assuming the role of 'expert' by explaining to other members of the group what the implications of data science might be. But they were far from homogeneous in their reactions to specific projects. For some, awareness of how data science works was coupled with optimism and openness towards government using data science to improve services. While to others, greater technical understanding led to greater scepticism and a belief that government use of data science involved a potential threat both to one's personal privacy and wider public benefit.

The research **challenges the notion that experience of sharing data breeds favourability towards data science**. As the latter case study shows, there is evidence to suggest the contrary and the fact that familiarity can bring with it a whole new set of questions and concerns that in turn sway an individual's risk assessment and make them less favourable.

When communicating with the public about future data science projects, it will therefore be important not to assume that the more the public know about how data science works, the more accepting they will be.

But if familiarity with data science is not the hook, **personal experience of (and access to) the benefits of data science is more compelling**. Showing individuals how government using data science can have a positive impact on their own experience and bring personal benefits to them or people they know could therefore be valuable in building public support.

7. Ethical framework and governance

7 Ethical framework and governance

This chapter explores participants' thoughts on what broad principles policy-makers should take into account when conducting data science projects.

Summary

On day two of the reconvened workshops, participants were tasked with developing priorities for people working in data science. The elements that they identified as important were generally well-aligned with the six principles of the draft ethical framework. While they were fairly clear on how they would expect data science projects to be run, they were less clear on the practicalities of how these principles would be implemented in a real-world setting. This exercise highlighted a number of areas where the draft ethical framework could benefit from further development and clarification, namely in areas of transparency, outcomes, intended consequences and the criteria for robust data science models.

7.1 Developing ethical principles

Participants were first of all tasked with developing a set of principles which they felt government should adhere to when carrying out data science projects. As part of this process they were asked to identify what they thought were the key priorities for people conducting these projects, including any considerations around privacy and safeguarding. Following this, participants were shown the six principles for good data science that GDS has designed as part of the draft ethical framework in order to establish their understanding of each of these principles and whether they felt they were important as far as conducting good data science is concerned.

Participants were broadly reassured that there was an existing ethical framework in place. In general, they talked in the same terms as the ethical framework. While the language used when conducting the original mapping exercise may not have been the same, the content and priorities echoed those detailed in the framework's six principles. When thinking about the rules that government should adhere to, they recognised that transparency, standards, data security, outcomes, benefits and privacy were all key considerations. These themes are explored in further detail below, including details of participants' ideas on each and how these can be mapped back to the principles of the draft framework. It is worth noting that the priorities participants came up with were very much focused on the general approach that the government should take in considering opportunities for data science and the overall approach, rather than specific principles that should be followed to help determine whether specific projects should go ahead.

7.2 Transparency

Transparency was a common theme throughout the workshops when discussing the principles of data science. Participants felt that this should be central at every stage of data science; as relevant to the data collection method, the security of that data, and the eventual outcome of the project itself. Being open and honest about why data is to be collected and for what purpose was important to participants. This can give them, and other members of the general public, reassurance around what the government is doing with their personal data.

Their intentions are better than Google, but they need to look after the people of Great Britain. You need to feel that you trust them.

Wolverhampton, 'high data' event 2

Raising public awareness of data science projects was deemed an important step in this process, creating an open and honest dialogue with the general public, in order to work towards a relationship of trust. Clearly stating the benefits of these projects, educating the public so that they have an understanding of the issues surrounding them and giving them contextual information around why, how and when their data is being used would go some way to increasing public confidence when it comes to data science. There were also suggestions that they would like to be given information post-project on the outcome or impact of data science projects.

I agree it's about being open and upfront about what the outcome of gathering the data has been – that's the bit we don't hear about! We don't know what changes have been made in Taunton as a result of the data being collected. None of us know the impact our data collection has had on the local community. As a taxpayer, how is that money being spent.

Taunton, event 1

I think they'd need public advertising and info campaign to convince the public that it's something that's worthwhile doing. Government want people to be on board with data science and to see it as something that's not scary, but that it is something that can be used and understood to produce benefits.

Sheffield, event 2

For others it was less about proactive communications and more about making the information available to those who want it. Broadly speaking they seemed to be keen on having access to information rather than having any kind of public consultation on a project by project basis.

You should be able to make FOI requests – e.g. what current sources of datasets are you using? It should be more publicised that people can make FOI requests.

Taunton, event 2

Participants distinguished between different types of data science projects. For example they felt that it was important for them to be informed if data was being collected for an end that would affect aspects of their daily lives and local communities. However, they also recognised that they would not be made aware of all

uses of data science, for example those that involve the police or security services. We explore the public benefits of data science projects in more detail later in this chapter.

The public voice wouldn't be asked whether they should investigate a terrorist, but it should be about say, should be' look at all this data about hospitals'. If they are looking at data which affects our lives, they [the public] should be consulted.

Taunton, event 2

This notion of transparency aligns well with principles 4 and 5 of the draft ethical framework: 'Be alert to public perceptions' and 'Be as open as possible'. When participants were shown these principles, they wondered about the practicalities of how the government could be alert to public perceptions. Would this, for example, require a public consultation each time a data science project was going to be conducted? It was important to them that the public voice would be heard in these situations, but they also understood that, in practical terms, minority claims cannot affect the overarching aims of a project if it is for the good of the population as a whole. When it comes to being as open as possible, participants broadly agreed that this was an important part of being transparent, however several queried how you would define the limits of being as open 'as possible' and whether this would always be appropriate under circumstances that would put lives at risk. This was therefore felt to be a somewhat loose definition that was open to interpretation.

Further discussion on engaging with the public on data science can be found in chapter 4.

It is also worth noting the potential research effect that taking part in the workshops could exert on participants' views on the need for transparency about data science. This discussion took part half way through the second day of workshops, so participants had already been presented with a lot of information about data science. Participants did not generally realise prior to the workshops that data science was being undertaken by government, or indeed that it is something that might be used increasingly by government as a means by which they make decisions that might impact on members of society. In turn this may have made them realise that most members of the general public are in fact not privy to this information as they are, hence the heightened sense that transparency is the way forward.

7.3 Standards and mechanics

Alongside transparency, participants wanted to know that data science projects would be run with integrity and impartiality, with high standards in governance, oversight and regulation paramount to their ultimate success and impact. They recognised that computers play an important role in data science projects in terms of sifting through vast amounts of data. However, they were hesitant about putting too much trust in computers and were emphatic about the role that humans should play in interpreting trends and correlations in the data and the final decisions made based on the outcome of data science projects.

People should be involved in decision making when the risk is related to human beings or businesses. Don't want to become too reliant on machines or 'blame' algorithms for decision making.

Taunton, event 2

To do this robustly – it must always be numbers and opinions combined.

Taunton, event 2

Participants saw governance and oversight as key priorities for people who work in government and who use data science. The integrity of data science projects was deemed important and participants felt that having a set of common standards which data scientists should adhere to is essential. These standards should be common values that are also imposed on any agencies or partners that the government involves in data science projects. These standards or criteria would be important in ensuring that data is collected appropriately, is not misinterpreted, and that findings are not over-generalised or used in a partial way against particular groups or sections of society. Ensuring the impartiality of data scientists should also be a key part of this accreditation process.

The workshop participants also recognised that the data sources should be reliable. Social media was felt to be less reliable than some other sources of data, however participants could still see the benefits of using this as a when looking at particular issues (citing the cast study on social media buzz ahead of the Norovirus outbreak). They also felt that any form of data should not stand alone but where possible be linked in with other sources to ensure that it has not simply been created in a bubble. It needs to be of sound quality to stand up to close inspection.

Quality of the data, accuracy of the data. Integrity of the process. You'd want to be able to see references in the reports, see where they got the data from. Make sure you're not manipulating the data to build the report.

Taunton, event 2

There were spontaneous mentions of the need for regulation and for a code of practice around data science. Participants felt that regulation should be carried out by an objective person or organisation and that this would add an extra layer of legitimacy and robustness to data science. There was the assumption among participants that if the government is using data science then it must be well regulated and be safeguarded, and that there should be a code of practice in place. Indeed, participants recognised that issues around safeguarding need to be considered at a number of stages in a data science project.

"I appreciate the need for a specialist but I think you need an independent moderator as well, because the scientist might not be thinking of the people behind the numbers."

Taunton, event 2

Concerns were expressed about what happens when regulation fails or doesn't exist for unforeseen issues, for example in cases of hacking or other misuses of data. Participants felt that it was important that there should be a complaint or error process in place to make sure that individuals had somewhere to go should their data had been misused, as an avenue to seek redress for mistakes or negative consequences. For some, the fact that they were asked about regulation during the workshop made them worry that this is not currently being as well-policed as it could be.

Closely related to this, participants discussed whether individuals should be able to opt-in or out of having their personal data included in data science projects, and whether the government or other people using data science should gain permission to gather data on individuals. There were mixed views about this, with some saying they were happy for the government to access their data, again based on their implicit trust that the government would safeguard their data in some way. Others were more guarded and felt data should only be gathered and used with express permission from individuals.

If it's personal and you can recognise a person then it's bad; but data security isn't important if you're not personal identifiable. Would be of less concern if it's a big dataset.

General public, Taunton event 2

Personal data should be kept personal – shouldn't be used by anyone other than the individual. Only used with consent or permission.

General public, Taunton event 2

While consent was clearly important to some participants, there was general acknowledgment that gaining permission to access personal data is not always practical. Participants were also able to spontaneously think of situations where consent would not be sought, for example when police are investigating crimes or incidences of terrorism.

The notion of having a set of standards by which people in government and working in data science adhere to links in with principle 3 of the draft ethical framework: 'Create robust data science models'. Members of the general public are clearly aware of the need to collect robust data and actually use this terminology themselves. However, there was less clarity around what a robust data science model would look like in practical terms. When presented with this principle from the framework, participants agreed that gathering reliable data was important; however they were unclear as to whether creating robust data science models is something that is achievable at this point in time, and also unclear as to how government would go about doing this.

A worthy ambition, but I expect tons of mistakes will emerge.

Wolverhampton, 'high data' event 2

Mistakes will always happen, but yes a system that evolves with the times [would be good].

Wolverhampton, 'high data' event 2

Reducing the risk of mistakes was the main priority, and participants certainly recognised that robust models should be in place where possible. In instances when errors do occur, indications were that publicly-acknowledging when and why these errors have occurred and making this information available to the public would go some way towards establishing public confidence and trust in data science and data scientists themselves.

While participants could identify that particular standards should be in place when it came to data science, they understood less well the mechanics of what made high quality, robust projects and therefore what content or

detail should go into the standards. They were however satisfied that those who do have such knowledge and understanding should set these standards.

As previously discussed, participants were less concerned about data being shared if the data was anonymised and it was being used for the greater good. In fact, some went as far as to say that data should be shared where possible, unless there was a good reason not to. This goes hand in hand with the assumption that data science is being used towards positive ends, something participants assumed would underlie any data science project.

The way in which participants prioritise issues around privacy mirrors closely principle 2 of the ethical framework: 'Use data and tools which have the minimum intrusion necessary'. When participants were shown this principle, they generally agreed that this was fair and that data scientists shouldn't go further than is absolutely necessary and should justify the level of intrusion used when conducting data science.

7.4 Data security

When thinking in more detail about the risks involved with data science, participants often assumed that this would contain personal information and quickly began discussing issues of data security and the safeguards that should be in place when it comes to protecting confidentiality. Participants were able to distinguish between different types of data – personally identifiable or anonymous – and felt that different rules should be applied accordingly. For example, many thought that government should take extra care and have stricter rules when it comes to processing and drawing conclusions from criminal justice data or data relating to individuals' sexual or mental health. There was also acknowledgement that armed forces and security services need to have stricter rules on data collection and retention.

More personal information needs better handling, take more care.

Sheffield, event 2

This prioritisation of data security by participants links into principle 6 of the draft ethical framework: 'Keep data secure', which stipulates that the government has a duty to protect the public's data. When participants were shown this principle it was generally met with comments such as '*obviously*', '*naturally*' and '*goes without saying*'. Some participants thought that keeping data secure was a given to the extent that they did not even think to mention it as a priority in the original exercise when identifying priorities for data science.

I didn't even think about keeping it secure. The obvious one too!

Taunton, event 2

The principle of keeping data secure was seen to be essential to the long-term success of data science and crucial for developing public confidence in the way government uses data science.

[There should be] as many safeguards as possible, not leaving data on trains or at least being totally accountable if mistakes, risks or leaks are made

Sheffield, event 2

Having strict rules on data security was seen as reducing the risk of deliberate or accidental data breaches. For some adhering to this principle would also be a good way of ensuring that the data was used for the purposes it was collected for, something that is discussed in more detail in the remaining sections of this chapter. Even so, others remained sceptical that their data would potentially be sold on to businesses for commercial purposes.

7.5 Outcomes and benefits

Participants had quite clear views on what a legitimate data science project would look like and how it would ideally be conducted. Their message is fairly clear when it comes to data collection: there should be a purpose and a proposed outcome. Without this in place they questioned why money would be spent on the project. While achieving outcomes was not a principle for data science as such, this was clearly a consideration for participants when it came to whether something should be approved or not.

Participants felt that a data science project should have clear objectives which should be both worthwhile and apolitical. They should be used for positive ends, for example making improvements to security, healthcare, education, and if these outcomes are tangible and they can see the immediate benefits in their local community, then all the better. Participants generally had strong feelings that the data should not be collected to then be sold on to private companies for gain, profit-making or for other purposes.

I personally would like to set politics aside, set objective aims, so they don't use it for campaigning, it has to stay within the civil service, it can't be used just as a cross bench argument. If it's going to be done it's got to be done by the state, not the government.

Taunton, event 2

Although a minority, some participants could see the benefits of re-purposing anonymous data if it was used to make improvements to public services or to have some other kind of positive social impact.

I think if it's not individual data, just general trends, I think repurposing is ok if for the greater good of society.

Taunton, event 2

The perceived power of the outcome can have a significant impact on how the general public establish whether the project is worthwhile. The workshop participants generally felt that the outcome should be proportionate to the risks involved in conducting data science in the first place. Generally speaking participants assumed there were public benefits in data science (see chapter 5 for further detail on this) and so did not dwell on this point at this stage of the discussion; however this is still an important consideration when it comes to the ethical framework.

As demonstrated in the case studies presented in section 5.2. The outcome of the project affected not only whether the project met test of 'public benefit' but also the mechanics of the potential project, and how much risk they were willing to take for both intended and unintended consequences. For those outcomes with the most at stake for a single individual or organisation (for example closing a food outlet, imprisonment, stopping

benefits) participants were less flexible in the type of data, methods used, and balance of false negatives to false positives.

The way in which participants talked about these issues around public benefits lined up with principle 1 of the ethical framework: 'Start with clear user need and public benefit'. When participants were shown principle 1 however, they felt that the terms 'clear user need' and 'public benefit' were slightly ambiguous and somewhat open to interpretation, and suggested these should be better and more clearly defined – who would the user be, how and who would define their needs, and who would the projects benefit, all or most of the public?

They were clear on the fact that the general public should be informed as to who is collecting the data, how the data would be used and why. The draft ethical framework discusses the risks of unintended consequences, such as how algorithms might use incomplete data or learn from historical bias, and that this might get in the way of the intended consequences or intended outcomes of data science. However it does not cover intended consequences in much detail. Given that outcomes have a clear impact on how the general public view data science, policy makers and data scientists using the framework would benefit from specific reference to how best to combine considerations of outcome, intervention and methodology together. This guidance could either form part of the assessment of 'public benefit' or be added as a new principle to the ethical framework.

7.6 Recommendations for the Ethical Framework

From the workshop discussions on priorities for people working in data science, it is clear that participants feel that transparency, security and safeguarding should be considered at all stages of a data science project. Improving awareness of data science projects among the general public would go some way to increasing the trust they have in what data of theirs is being collected and what it is being used for. During the discussions participants often came back to the intended outcomes of the data science projects, generally keen to know that there was a legitimate reason behind these projects being undertaken in the first place. They were most positive about data science projects that would have a clear positive outcome in their communities, resulting in, for example, improvements in healthcare or education. Providing the outcome was proportionate to the risks involved participants were broadly positive about data science as a process and its potential for having a positive impact on society.

On the whole, participants agreed that the six principles of the draft ethical framework covered some of the key elements of what they thought should be priorities for people working in data science. The workshops however brought to light some specific areas of the draft ethical framework that could benefit from further clarification or elaboration.

- **Transparency** – it would be beneficial to have more guidance on how people working in data science go about being transparent and communicating data science to the public as part of open policy making. Building on the current framework, best practice on gaining public trust and confidence would be helpful.

- **Outcomes** – the risk and proportionality of ‘outcomes’, both intended and unintended needs to be clearly recognised in more detail within the ethical framework, especially given the impact that this variable appears to have on the public assessment of whether data science opportunities are suitable.
- **Robust data science models** – further clarity around the composite parts of how a robust data science models is defined, what they look like and what standards they are subject to would be beneficial, especially when thinking about how best to be transparent and communicate data science to the public.

Steven Ginnis

Associate Director
steven.ginnis@ipsos.com

Harry Evans

Senior Research Executive
harry.evans@ipsos.com

Naomi Boal

research Manager
naomi.boal@ipsos.com

Emily Davies

Research Manager
Emily.davies@ipsos.com

Alexandra Palmqvist Aslaksen

Research Executive
alexandra.aslaksen@ipsos.com

For more information

3 Thomas More Square
London
E1W 1YW

t: +44 (0)20 3059 5000

www.ipsos-mori.com

<http://twitter.com/IpsosMORI>

About Ipsos MORI's Social Research Institute

The Social Research Institute works closely with national governments, local public services and the not-for-profit sector. Its c.200 research staff focus on public service and policy issues. Each has expertise in a particular part of the public sector, ensuring we have a detailed understanding of specific sectors and policy challenges. This, combined with our methodological and communications expertise, helps ensure that our research makes a difference for decision makers and communities.